Consensus Simulator for Organisational Structures

Johannes S. Vorster[1][a] and Louise Leenen[1,2][b]

[1]*Department of Computer Science, University of Western Cape, South Africa*

[2]*CAIR, South Africa*

Keywords: Consensus, Consensus Simulation, Stochastic Simulation, Synchronization.

Abstract: In this paper we present a new simulator to investigate consensus within organisations, based on organisational structure, team dynamics, and artefacts. We model agents who can interact with each other and with artefacts, as well as the mathematical models that govern agent behaviour. We show that for a fixed problem size, there is a maximum time within which all agents will reach consensus, independent of number of agents. We present the results from simulating wide ranges of problem sizes and agent group sizes and report on two significant statistics; the time to reach consensus and the effort to reach consensus. The time to reach consensus has implications for project delivery timelines, and the effort relates to project economics.

1 INTRODUCTION

The past two decades have seen debates on shifts in organisational structures and project delivery methodologies. These debates originated with technology companies that needed to cope with three factors, namely, changes in technology, changes in competitor position, and fast shifting customer demands. Often quoted examples of such changes are shifts in team structure (Reagans et al., 2016), shifts from vertical to more horizontal organisations (Keupp et al., 2012), and organisational structures that resemble network-like forms (Chang and Harrington, 2000). The outcomes seem to be leaning towards the conclusion that more horizontal organisational structures and better connected open-network structures create better economic value in the form of faster delivery and less resources (time and material) spent on projects and thus deliver better returns on investments (Will et al., 2019). Project complexity and the ability to find consensus on approaches and solutions have been identified and studied as key reasons why large projects fail (Kian et al., 2016).

Will et al. (2019) studied how organisational structure affect the economics of accepting risky projects. In particular, they study the effect of the selection process and the economic impact of organisational structure on risky and innovation projects selection. They argue that the selection of innovations to pursue from an available portfolio is not only dependent on the team's evaluation skills, but also on the organisational structure. They approach the topic with a purely mathematical model. They showed that, for example, hybrid organisational architectures tend to have side effects in terms of handling errors. Sáenz-Royo and Lozano-Rojo (2023) conducted similar structures to investigate innovation project selection.

Human consensus in formal settings, such as prediction of economic outcomes, are often conducted by using Delphi processes where a group is led to consensus though repeated rounds of providing views anonymously. There are other methods to study consensus, qualitative studies of social networks and consensus emergence within these networks (Carter et al., 2015; Jones and Shah, 2016). A second approach is to conduct detailed interviews and study the phenomenon qualitatively (Roselló et al., 2010). A third approach is to study consensus through computational models (Yan et al., 2017). This paper employs the last approach, but also incorporates aspects of Delphi decision making.

In this paper we report on the implementation of a simulator for the study of organisational structures, team dynamics and economic implications. We observe that project implementations are a series of consensus-seeking processes, where teams must agree on vision, scope, requirements, architecture, design, implementation details, construction decisions, security controls, testing, quality control measures, and so on.

[a] https://orcid.org/0000-0001-6452-4186

[b] https://orcid.org/0000-0002-9212-550X

Vorster, J. and Leenen, L.
Consensus Simulator for Organisational Structures.
DOI: 10.5220/0012017400003546
In *Proceedings of the 13th International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH 2023)*, pages 15-26
ISBN: 978-989-758-668-2; ISSN: 2184-2841

We follow a stochastic agent-based modelling approach, where connections between agents simulate organisational structure and communications channels as well as access to artefacts. Agents have prior opinions about topics, which can be modelled using various statistical distributions. The model simulates the best-case consensus effort and time for the configuration of organisational architecture and team structure to reach consensus.

An aspect important to this simulation is that agents manage their time by keeping diaries and cannot have multiple meetings at the same time. Meetings are the means to settle differences in views and agents are restricted in the number of topics that can be discussed per meeting.

We provide details for the model and the simulator, as well as initial results regarding the impact of artefacts on consensus seeking processes and organisational efficiency.

In the next section the model is described, which is followed in section three with a mathematical treatment of the model. In section four the results from simulating various scenarios are presented and discussed.

2 MODEL DESCRIPTION

In this section we discuss the construction of our model of consensus and the simulator that was built to implement that model. Topics covered are organisational architecture, the graph theory model, consensus measures, artefacts, agent actions, diaries, and effort & time measures.

2.1 Variations in Organisational Architecture

The two extreme forms of organisational architectures are Polyarchies (Figure 1) and Hierarchies (Figure 2) (see Will et al. (2019) for a discussion). Hybrid organisations (Figure 3) are somewhere between these two architectures.

2.1.1 Polyarchical Organisations

In polyarchical organisations (Figure 1) the agents are peers with the same capabilities, organisational power, and connectedness. Agents have a view of every other agent's position on topics and they can take actions to reach consensus with each other. If agents share artefacts such as plans or requirements, they can simultaneously work on parts of the artefacts.

2.1.2 Hierarchical Organisations

In the case of a hierarchical organisation (Figure 2), managers on different levels also differ in decision-making powers. Managers on the same level have different executive functions. In this organisational structure, information flow and consensus-seeking processes must follow the hierarchical rules in that subordinates convey information only up and down the 'chain of command' and thus this hierarchy also places constraints on the way consensus-seeking tasks such as meetings are conducted. An agent can only organise a meeting with peers, subordinates or their manager.

In such an organisation, the top-level managers may provide a project vision, which is translated into a requirements specification by one group, an architecture by a second group, and a detailed delivery plan by a third group. The artefacts that are finally delivered (the software system in the case of such a project, or the road system if that was the project) can be measured against the expectation (the vision) of the top-level manager, and against other planning artefacts.

2.1.3 Hybrid Organisations

The hybrid organisational architecture (Figure 3) is a combination of polyarchical and hierarchical structures (Christensen and Knudsen, 2010; Will et al., 2019). The organisation has a hierarchy, but agents are not forced to communicate and structure decisions within that hierarchy and can communicate with peers in other teams or departments. This facilitates consensus processes to work both vertically and horizontally (Young-Hyman, 2017).

2.2 Graph Theory Preliminaries

Agents will be denoted ν_i for the i^{th} agent from a set of N agents. In classical agent consensus theory the agents are represented as vertices in a graph and the edges represent the connections between agents (e.g. Wei et al. (2021)). We extend those definitions to include artefacts.

Let $G = (V, \varepsilon, \Delta)$ be a directed graph, where $V = V_\nu \cup V_A$ the set of vertices with $V_\nu = \{\nu_i | i \in I_\nu = \{1, 2, \ldots, N\}\}$ the vertices that represent N agents and $V_A = \{c_p | p \in I_A = \{1+N, \ldots, M+N\}\}$ the vertices that represent M artefacts. The index set over V is then $I = I_\nu \cup I_A$.

The set of edges is represented by $\varepsilon \subseteq V \times V$. Then $e_{ij} \in \varepsilon$, $i, j \in I$, represent an edge of G that correspond to an interaction. The element δ_{ij} in the adjacency matrix $\Delta = [\delta_{ij}]$ correspond to e_{ij}, and δ_{ij} is positive if and only if $e_{ij} \in \varepsilon$, otherwise $\delta_{ij} = 0$.

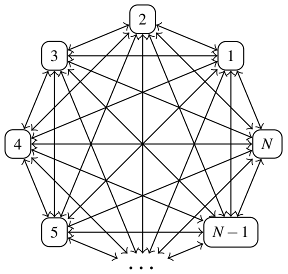

Figure 1: Polyarchy.

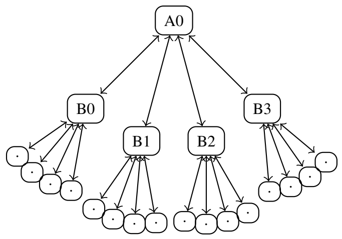

Figure 2: Hierarchical.

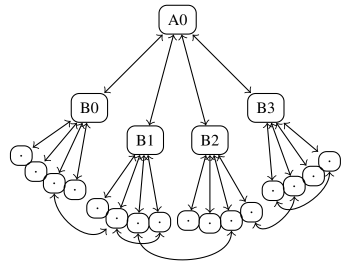

Figure 3: Hybrid.

If $\nu_i, \nu_j \in V$ then information can flow from ν_i to ν_j if $\delta_{ij} > 0$. A graph is directed if, for some $i, j \in I$ $\delta_{ij} > 0$ but $\delta_{ji} = 0$. That is, ν_i can interact with ν_j, but ν_j cannot interact with ν_i. Otherwise it is undirected. Artefacts cannot interact with each other and thus $\delta_{pq} = 0, \forall p, q \in I_c$.

The numeric value δ_{ij} can represent some characteristic of the interaction such as the amount of information transmitted, the bandwidth, or the probability of successful communication. In the work presented here

$$\delta_{ij} = \begin{cases} 1 & e_{ij} \in \varepsilon_a \\ 0 & \text{otherwise.} \end{cases}$$

2.3 Topics, Views, and Artefacts

An agent $\nu_i \in V_\nu$ can have a view (synonymous to opinion and belief) on a number of topics. If k is an index number to the topics, $k \in I_B = \{1, 2, \ldots, \mathcal{B}^{\max}\}$, then b_i^k denote the view of agent i on topic k. Topics are independent of each other since each topic is about something specific. Topic k is the same topic for all agents, although each agent can have a different view on that topic. Thus b_i^k, agent i's view on topic k, can be compared with b_j^k, agent j's view on topic k. However, b_i^k cannot be compared with b_j^m, if $k \neq m$.

An artefact represents a set of codified (written down, or captured in some other way) views on topics. Topic k is codified in artefact p by c_p^k. There may be many artefacts c_p, $p \in I_c$ and many topics $k \in I_C = \{1, \ldots, \mathcal{C}^{\max}\}$. Since c_p^k is the codification of topic k, it is possible to compare c_p^k with b_i^k, $k \leq \min(\mathcal{B}^{\max}, \mathcal{C}^{\max})$. Note that $\mathcal{C}^{\max}$ and $\mathcal{B}^{\max}$ could be different. That is, the artefact could cover more topics ($\mathcal{C}^{\max} > \mathcal{B}^{\max}$), or less, or the same number of topics than what agents may have views about.

2.4 Modelling Consensus

Next we address the question of how to compare topics with each other. We are interested in a measure that reflect an agents' level of consensus with the other agents in its network as well as an overall measure of group consensus.

In agent-based modelling, consensus between an agent and the rest of its connected group on a specific topic k is often expressed as the sum of differences between agent i and all other agents j, (see e.g. Wei et al. (2021)),

$$u_i^k = -\sum_{j=1}^{N} \delta_{ij}(b_i^k - b_j^k), \quad i \in \{1, 2, \ldots, N\}.$$

This definition computes the average difference, and can be negative. This way of measuring consensus has the aim to indicate to the agent in what direction (positive or negative) the group average is located. The agents will then modify their view to move towards the group average.

However, this definition of consensus does not suite us here, since it leads to the strange situation where if $b_1^k = 0$, $b_2^k = 5$, and $b_3^k = 10$, then $u_2^k = (b_2^k - b_1^k) + (b_2^k - b_3^k) = 0$ which seem to imply that it is in consensus with the other agents whereas it only holds the average view.

In human consensus processes, the most famous of which is the Delphi process, consensus can be measured in may ways, one of which is the sum of a pairwise comparison of views between agents (see e.g. Birko et al. (2015) for the many ways to measure consensus in Delphi processes).

We are interested in how far away the group of agents are from a state of consensus, and thus opt for the sum of absolute differences as the consensus measure, and sum over all pairwise agents and all topics. A full justification and comparison with other alternatives is beyond the scope of this paper. We thus opt for a measure that is a combination of consensus measures from Delphi (pairwise comparison of consensus) and agent-based modelling. We use absolute differences since we want to capture how far agents are from consensus. Furthermore, we include artefacts in the evaluation and measure how well agents are in consensus with the artefacts.

We used the approach of defining a measure of consensus for agents with each other ($i, j \in I_\nu$) and agents with artefacts ($i \in I_\nu, p \in I_A$) as

$$u_{ij}^k = \delta_{ij}|b_i^k - b_j^k| \text{ and } u_{ip}^k = \delta_{ip}|b_i^k - c_p^k|$$

which leads to an overall measure of consensus for an agent i with agents and artefacts it has contact with as

$$u_i^k = \sum_{j \in I_\nu} \delta_{ij}|b_i^k - b_j^k| + \sum_{p \in I_A} \delta_{ip}|b_i^k - c_p^k|. \quad (1)$$

That is, the level of consensus that an agent i has relative to the rest of the group on a topic k, is the sum of absolute differences between that view b_i^k and the views on the same topic for all other agents, b_j^k that it is connected to, ($\delta_{ij} > 0$), as well as the same measure for that topic in all artefacts, c_p^k, that it has access to.

The consensus of an agent with its group of connected agents and artefacts over all topics is given by

$$u_i = \sum_{j \in I_\nu} \sum_{k=1}^{\mathcal{B}^{\max}} \delta_{ij}|b_i^k - b_j^k| + \sum_{p \in I_A} \sum_{k=1}^{MX} \delta_{ip}|b_i^k - c_p^k|, \quad (2)$$

where $MX = min(\mathcal{B}^{\max}, \mathcal{C}^{\max})$, see section 2.3.

The overall consensus on a specific topic k can also be defined using (1) as

$$u^k = \sum_{i \in I_\nu} \sum_{j \in I_\nu} \delta_{ij}|b_i^k - b_j^k| + \sum_{i \in I_\nu} \sum_{p \in I_A} \delta_{ip}|b_i^k - c_p^k|. \quad (3)$$

Finally, we can now define an overall level of consensus over all agents, that is, an overall level of consensus for the group as

$$u = \sum_{i \in I_\nu} u_i. \quad (4)$$

As will be shown shortly in section 3.3, the consensus follows an exponential decrease. Therefore, we also represent the consensus as the natural log of u

$$S_u = \ln(u). \quad (5)$$

This measure of consensus is more descriptive and insightful than using u, as will be discussed in the results section, 4.2.

We want to be able to model the situation where agents may believe that they are in consensus with another agent or artefact, but in fact are not. This is a common situation, where people believe they are in agreement, but subsequent more detailed analysis shows that they are in fact not in agreement. This is achieved by introducing a consensus threshold, κ. If the absolute difference in views on a topic is within this threshold, the agents will consider that topic as synchronized and no further discussions or actions will involve that topic. Consensus is reached when $u_{ij}^k = \delta_{ij}|b_i^k - b_j^k| \leq \kappa$. That is, the agents will believe that consensus is not reached if $u_{ij}^k > \kappa$ and will take further actions to achieve consensus on topic k with agent j. The same holds for artefact consensus and if $u_{ip}^k = |b_i^k - c_p^k| \leq \kappa$ the agent will consider that it is in consensus with artefact p on topic k.

Consider an example of the consequence of the introduction of κ. The project sponsor generates a project vision document where his actual vision differs from what is documented by κ. The business analyst reads this document, again with a κ margin, which now creates a $2 \cdot \kappa$ difference between what the analyst understands and what the project sponsor meant. In this way, a chain of agents and artefacts can create an ever increasing 'error' between the original meaning and what is understood later in the chain.

In all the simulations conducted and reported on later $\kappa = 2$.

2.5 Agent Actions and Diaries

Artefacts could be assigned to every member of the group. The assignment also specifies the actions that each agent is allowed to take on its assigned artefact, namely writing to the artefact, reading from the artefact, or both. One group of agents may be responsible for writing an artefact, for example project requirements, software source code, or the architecture document for a new road network. While another group may only read that artefact, for example the construction team can only read the architecture document and cannot directly change it.

Agents can take only one action in a time interval. At the beginning of the time interval, each agent (ν_i) constructs a plan which consists of all potential actions available. This plan consists of agents and artefacts that ν_i has connections to, and that have topics where consensus has not been reached. Therefore the agent will identify all ν_j, $\delta_{ij} > 0$, for which there exists a k so that $|b_i^k - b_j^k| > \kappa$. Similarly they will identify all artefacts c_p, $\delta_{ip} > 0$ for which there exists a k so that $|b_i^k - c_p^k| > \kappa$.

An agent selects at random from the set of actions in its plan. These actions are:

- Meet with another agent, ν_j. A meeting is recorded in both agents' diaries. A random number of topics, d, are discussed, where $|b_i^k - b_j^k| > \kappa$, $d \in I_D = \{1, \ldots, \mathcal{D}^{\max} - 1\}$ with expected value $\hat{d} = \mathcal{D}^{\max}/2$. For every topic, k, discussed, one of three outcomes is selected at random; (a) a compromised consensus where the new b_i^k and b_j^k are set to the average of their values at the beginning of the meeting, (b) ν_i accepts the view of ν_j by setting $b_i^k = b_j^k$; or (c) ν_j accepts the view of ν_i by setting $b_j^k = b_i^k$.
- Read an artefact p if an artefact is present and

the read action is allowed. The agent interacts with the artefact and updates its own views on a random number, $d \in I_D$, of topics where consensus has not been reached, by setting $b_i^k = c_p^k$, $k \leq min(\mathcal{B}^{\max}, \mathcal{C}^{\max})$.

- Write to artefact p if an artefact is present and the write action is allowed. Here the agent imparts its views into the artefact, setting $c_p^k = b_i^k$ for a random number of topics, $d \in i_D$ and $k \leq min(\mathcal{B}^{\max}, \mathcal{C}^{\max})$.
- Do nothing. If an agent cannot construct a plan, because it is in consensus with all agents in its network and all artefacts it has access to, then it will do nothing during that time interval.

The order in which agents take action is randomized at the beginning of each time interval. If an agent (ν_j) was involved in a meeting with another agent (ν_i) due to ν_i taking an action earlier, then ν_j was already involved in an action this time interval and cannot take a further action.

An agent's diary captures the action it took in that time interval. We use d_i^t to represent the diary action of ν_i at time t. It is a historic record that can be analysed later to understand how productive an agent was and other behavioural statistics. If ν_i was the originator of a meeting, it is indicated with $d_i^t =$ 'O'. If a meeting was attended, but ν_i was not the originator, then $d_i^t =$ 'm'. For the actions to read, write, or update an artefact the diary entries are $d_i^t =$ 'r', $d_i^t =$ 'w', and $d_i^t =$ 'u' respectively. If the agent could not take any action then $d_i^t =$ 'z'.

2.6 Effort, Time, and Productiveness

The effort, $e^{\max}$, to reach consensus is the sum of all actions taken by all agents, that is

$$e^{\max} = \sum_{t=1}^{t^{max}} \sum_{i=1}^{N} busy(d_i^t),\ busy(d_i^t) = \begin{cases} 0 & d_i^t = \text{'z'} \\ 1 & \text{otherwise,} \end{cases} \tag{6}$$

where $t^{\max}$ is the total time it took to reach consensus. Since each agent will always take some action if an action is available, the simulation terminates at time t^{max} when $d_i^t =$ 'z' $\forall i \in I_\nu$. That is, when all agents no longer takes any actions, the simulation stops. Both $e^{\max}$ and $t^{\max}$ will be determined by simulation and will differ on each stochastic simulation.

These two variables, i.e. the effort measured by $e^{\max}$ and time by $t^{\max}$, are the measures we are most interested in. This paper started off by asserting that all projects are consensus-seeking processes that can complete successfully only if consensus is reached through the various phases and deliverables of the project. The *effort* can be translated into an economic measure by noting that it measures agents' time spent on the project, and time multiplied by rate gives cost. The overall *time* to complete the project (that is, calendar time) also has economic implications if opportunity costs and Net Present Values are considered.

In subsequent work we want to establish the characteristics of these two variables as a function of the organisational structure, project delivery methodology and even team dynamics. However, here we focus on the simulator and the basic results using polyarchy organisations as a benchmark for future work.

The configurations discussed in this paper is such that the agents are almost fully engaged and productive. That is not the case in all configurations. Part of what makes this interesting is that agents may become very unproductive for large periods in specific configurations. This may be the case where a subset of agents are responsible for writing requirements and another team is responsible for the implementation of those requirements. Initially the requirements writing agents will be engaged in artefact writing activities, but then they may be idle for a while while other agents assimilate the requirements, followed by a frenzy of activities as the agents try to resolve differences of views and reach consensus. These configurations are not reported here, but form part of the long term initiative we aim to report on in future papers.

3 MATHEMATICAL MODEL PRELIMINARIES

This paper does not expand on the full mathematical model, since the main aim is to model complex project configurations that may involve complexities better handled with simulation. However, some of the observed trends for polyarchies can be explained with mathematical models and so we offer a non-rigorous formalism here. Furthermore, if the simulation adheres to the model, it acts as validation that the implementation performs as expected and thus were implemented correctly. This bolsters confidence in the results produced by later simulations with more complex configurations.

This section assumes no artefacts are present. A group of N fully connected agents in a polyarchy will be called an N-group. For example a group of 15 fully connected agents will be called a 15-group.

3.1 Nomenclature Revisited

Agents (index set I_ν) in an N-group has opinions about $\mathcal{B}^{\max}$ topics (index set I_B). An artefact contains

$\mathcal{C}^{\max}$ codified topics (index set I_C). The maximum number of topics that can be discussed and concluded per action is $(\mathcal{D}^{\max} - 1)$, the minimum is one, the expected number (denoted with a hat) is

$$\hat{d} = \mathcal{D}^{\max}/2. \tag{7}$$

The topics, a total of $\mathcal{B}^{\max}$, is fixed for a given simulation and all agents within the group has $\mathcal{B}^{\max}$ topics

3.2 Initial Conditions

If $f(x,t)$ is the distribution function for the agent views (b_i^k) at time t, then the *expected value* of the mean absolute difference of the consensus measure between the views of two agents $|b_i^k - b_j^k| = u_{ij}^k$ is given by

$$\hat{u}_{ij}^k(t) = \int_{x=-\infty}^{\infty}\int_{y=-\infty}^{\infty} |x\cdot f(x,t) - y\cdot f(y,t)| \tag{8}$$

and thus, depending on the distribution,

$$\hat{u}_{ij}^k(t) = \begin{cases} \frac{1}{3}\Delta_b & \text{Uniform}(\Delta_b) \\ \frac{2\cdot\sigma}{\sqrt{\pi}} & \text{Normal}(\mu,\sigma) \\ \frac{1}{\lambda} & \text{Exponential}(\lambda) \\ \frac{3}{2\lambda} & \text{Symmetrical Exponential}(\lambda) \end{cases} \tag{9}$$

on topic k. Real world distributions of opinions have been shown to have Uniform (Den Boon and Van Meurs, 1991), Normal (Den Boon and Van Meurs, 1991) and Exponential (Lang et al., 2018) distributions, depending on the specifics of the topic. We make the assumption that opinions are Normally distributed for some topic so that $\hat{u}_{ij}^k = \frac{2\sigma(t)}{\sqrt{\pi}}$ and $\sigma(t)$ is time-dependent since the distribution will change over time.

At t=0 all agents are initialized with random views per topic, $b_i^k \sim \mathcal{N}(\sigma = 100)$, from the Normal distribution.

The expected value for $u_{ij}^k(t=0)$ in our simulations with $\sigma(0) = 100$ is then

$$\hat{u}_{ij}^k(0) = \frac{200}{\sqrt{\pi}}. \tag{10}$$

The group consensus measure, given $\hat{u}_{ij}^k(t)$, using (4), is

$$\hat{u}(t) = \sum_{i\in I_\nu}\sum_{j\in I_\nu}\sum_{k=1}^{\mathcal{B}^{\max}} \hat{u}_{ij}^k(t) = \mathcal{B}^{\max}N(N-1)\hat{u}_{ij}^k(t) \tag{11}$$

keeping in mind that $u_{ii}^k = 0$. At $t = 0$

$$\hat{u}(0) = \sum_{i\in I_\nu}\sum_{k=1}^{\mathcal{B}^{\max}} \hat{u}_i^k = \mathcal{B}^{\max}N(N-1)\frac{2\sigma(0)}{\sqrt{\pi}}. \tag{12}$$

The simulation ends when $|b_i^k - b_j^k| < \kappa, \forall i,j$ at $t = t^{\max}$. The expected value of $|b_i^k - b_j^k|$ is then $\kappa/2$ and thus

$$\hat{u}(t^{\max}) \approx \frac{1}{2}\kappa\mathcal{B}^{\max}N(N-1). \tag{13}$$

3.3 Consensus Change

Agents decide randomly between one of three options to resolve a topic with another agent (or an artefact) as discussed in section 2.5. For option (a) both compromise to the mean value of their views, in option (b) ν_i adopts the position of ν_j and option (c) is the reverse. Each of these events happen with 1/3 chance and the expected change in consensus due to each of these events can be computed using expected values from the Folded Normal distribution and summing using (4). The derivation is too lengthy to replicate here and we only provide the result by case. For option (a)[1]

$$\Delta\hat{u}_{ij}^k(t)_{(a)} = -\frac{1}{3}\hat{u}_{ij}^k\cdot 2\left[(2\sqrt{3}-3) + (2-\sqrt{3})N\right] \tag{14}$$

for cases (b) and (c)

$$\Delta\hat{u}_{ij}^k(t)_{(b,c)} = -\frac{1}{3}\hat{u}_{ij}^k(t)\cdot 2 \tag{15}$$

so that for all three cases

$$\begin{aligned}\Delta\hat{u}_{ij}^k(t) &= -\frac{2}{3}\left[(2\sqrt{3}-1) + (2-\sqrt{3})N\right]\hat{u}_{ij}^k(t) \\ &= -g'(N)\cdot\hat{u}_{ij}^k(t)\end{aligned} \tag{16}$$

where

$$g'(N) = \frac{2}{3}\left[(2\sqrt{3}-1) + (2-\sqrt{3})N\right] \tag{17}$$

This result is dependent on the choice of initial distribution for b_i^k, and consequently $\hat{u}_{ij}^k$ from (9). Figure 4 shows the distribution of $u_{ij}^k = |b_i^k - b_j^k|$ as well as its evolution for various time instances, collected over many simulations. These graphs uses a log-y scale, $b_i^k(t=0) \sim \mathcal{N}(\sigma = 100)$, and thus $u_{ij}^k(t=0)$ has a Folded Normal distribution. The evolution of the distribution of u_{ij}^k causes the model presented by (16) to become less accurate in later time

[1] Breadcrumbs to verify option (a) results: $b_i^k \sim \sigma_i$, $b_i^k - b_j^k \sim \sigma_{i-j} = \sqrt{2}\sigma_i$, $u_{ij}^k = |b_i^k - b_j^k|$ is Folded-Normal, thus $u_{ij}^k \sim \sqrt{\frac{2}{\pi}}\cdot\sigma_{i-j} = 2\sigma_i/\sqrt{\pi}$. Option (a): $m = \frac{b_i^k + b_j^k}{2}$, $\sigma_m = \sigma_{(i+j)/2} = \sigma_i/\sqrt{2}$, $\sigma_{m-i} = \sqrt{\sigma_m^2 + \sigma_i^2} = \sqrt{3/2}\sigma_i$, $|m - b_i^k| \sim \sqrt{2/\pi}\sqrt{3/2}\sigma_i$. Now sum $\Delta u^k = u^k(t+1) - u^k(t)$ using (4).

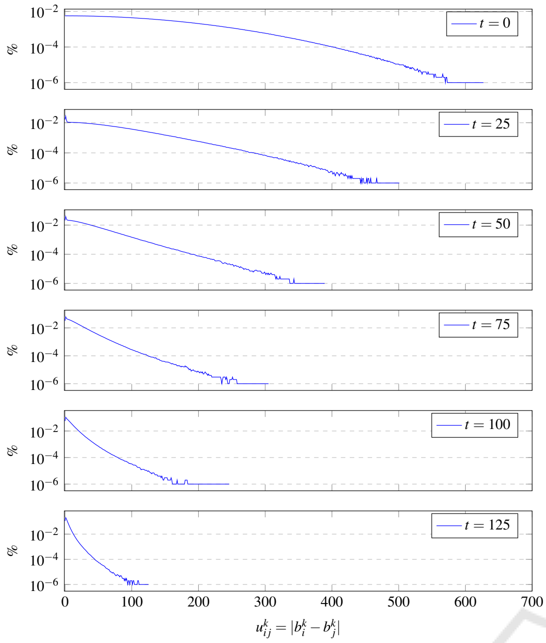

Figure 4: As the consensus process continues, the distribution of differences in opinions, $u_{ij}^k = |b_i^k - b_j^k|$, changes. Data shown are from simulations with a 50-group, $\mathcal{B}^{\max} = 50$ views on which to reach consensus, and $b_i^k(0) \sim \mathcal{N}(\sigma = 100)$.

steps. We do not see an easy way to build a model that tracks this change in distribution. One approach would be to introduce a sigmoid function to help correct for the change in distribution characteristics and we leave that for future work.

The expected change in the consensus measure for a single time interval at some time t, $\Delta t = 1$, is then the sum over all events, that is, through $N/2$ meetings wherein an expected $\hat{d}$ topics are discussed per meeting, leading to

$$\begin{aligned}\Delta\hat{u}(t) &= \sum_1^{N/2}\sum_1^{\hat{d}} \Delta\hat{u}_{ij}^k(t)\\ &= -\hat{d}\cdot\frac{N}{2}\cdot g'(N)\cdot\hat{u}_{ij}^k(t).\end{aligned} \tag{18}$$

Substituting $\hat{u}_{ij}^k(t)$ from (11) gives

$$\begin{aligned}\Delta\hat{u}(t) &= -\frac{\hat{d}\cdot N\cdot g'(N)}{2\cdot\mathcal{B}^{\max}N(N-1)}\cdot\hat{u}(t)\\ &= -\frac{\hat{d}}{2\mathcal{B}^{\max}}\frac{g'(N)}{N-1}\cdot\hat{u}(t)\end{aligned}$$

which can be written as

$$\begin{aligned}\Delta\hat{u}(t) &= -a_0\cdot\hat{u}(t),\\ a_0 &= \frac{\hat{d}}{2\mathcal{B}^{\max}}\cdot\frac{g'(N)}{N-1}\\ &= \frac{\hat{d}}{3\mathcal{B}^{\max}}\cdot\frac{(2\sqrt{3}-1)+(2-\sqrt{3})N}{N-1}\end{aligned} \tag{19}$$

where a_0 is now dependent on $\mathcal{B}^{\max}$, $\hat{d}$, and N, but static for a simulation.

3.4 Consensus

The expected value for u , at $t+1$ is $u(t+1) = \hat{u}(t) + \Delta\hat{u}(t)$, which we can now compute by using $\hat{u}(t)$ from (11) and $\Delta\hat{u}(t)$ from (19) to give

$$\hat{u}(t+1) = (1-a_0)\hat{u}(t)$$

and thus

$$\hat{u}(t) = \hat{u}(0)\cdot(1-a_0)^t. \tag{20}$$

3.5 Meeting Assumptions

The above equations that describe the measure of consensus and consensus delta per time interval are subject to the assumption that $\mathcal{D}^{\max}$ topics are available for discussion. However, as the simulation progresses, a time is reached when agents have reached a level of consensus with other agents so that the number of topics available between an agent i and j is less than $\mathcal{D}^{\max}$.

At that time, the above equation is no longer valid - at least in the sense that the number of topics that can be discussed per meeting is no longer $d \in I_D$, since there are less than $\mathcal{D}^{\max}$ available topics to discuss. That is, the expected number of topics discussed per meeting, $\hat{d}$ starts to decrease, meetings are less efficient, and $\Delta\hat{u}(t)$ decreases as more and more topics reach the consensus threshold κ.

As an example, it may be that agent i has five topics to reach consensus on, but, each one is with a different agent. In that case the agent still needs to attend five meetings one with each of the agents, but each meeting can only discuss one topic, making the meeting itself inefficient. This is not uncommon in real life.

The situation is illustrated in Figure 5, which shows the consensus trajectory (averaged over n=25000 simulations) for a 20- and a 30-group. This figure shows the model prediction using (20) and $\hat{d} = \mathcal{D}^{\max}/2$.

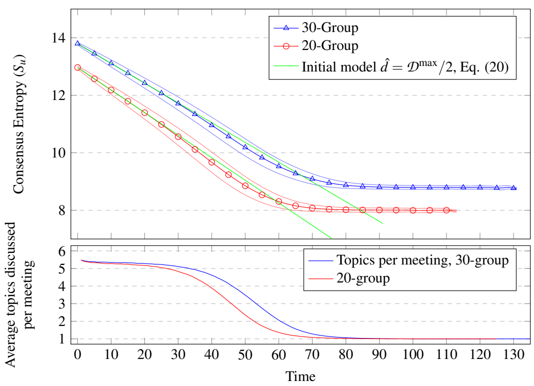

Figure 5: Later in the consensus process when the number of topics that needs to be discussed between two agents are less than the expectation, $\hat{d} = 5.5$, the meetings become less efficient. The consensus process is dominated by one-topic meetings towards the end of the consensus process.

Two behavioural changes can be seen in Figure 5. Firstly, the rate of change increases (graph drops below predictive model) as the distribution of b_i^k slowly changes away from Normal, as shown in Figure 4. Secondly, as the number of topics available to discuss starts to drop below $\mathcal{D}^{max}$ and thus $\hat{d}$ and $\Delta\hat{u}(t)$ diminishes, the rate decreases and the graph turns almost horizontal as meetings become less efficient. Figure 5 (bottom) shows that towards the end the consensus process is dominated by one-topic meetings.

This is an interesting result since it provides further insights into why project projected timelines are often exceeded. The one-topics region is a significant portion of overall project time and thus warrants further study both theoretically, quantitatively and qualitatively.

3.6 Estimations

Due to this change in the effectiveness of meetings as agents come closer to reaching consensus, it is difficult to find good estimators for t^{max} and e^{max}. An approximation can be found by using (13) to estimate $\hat{u}$ at t^{max}, substitute into (20) and solve for t_1 as the first estimation of t^{max}. Thus

$$\hat{u}(0)\cdot(1-a_0)^{t_1} = \frac{1}{2}\kappa\mathcal{B}^{max}N(N-1)$$

which leads to

$$t_1 \approx \frac{\ln\left(\frac{\sqrt{\pi\alpha\kappa}}{4\sigma}\right)}{\ln(1-a_0)},\ \alpha = 1, \tag{21}$$

where the subscript of t_1 indicates that this is a first approximation, and α is a scaling factor. If we want to compute t_1 to be the expected value at consensus, then $\alpha = 1$. We will shortly use α as a scaling factor to indicate when the consensus process reaches a time of low meeting efficiency, where there are much less than $\hat{d}$ topics available for discussion.

As a second approximation, we compute t to reach $\alpha\kappa$ for every topic and then do a second estimation of time to reach the expected $\kappa/2$ by using $\hat{d} = 0.2$. Use the first approximation, t_1, to estimate $u(t_1)$, and finally solve for t_2, the second estimate, in

$$\begin{aligned}\hat{u}(t_2') &\approx \hat{u}(t_1)\cdot(1-a_1)^{t_2'}\\ &= \hat{u}(0)\cdot(1-a_0)^{t_1}\cdot(1-a_1)^{t_2'}\end{aligned} \tag{22}$$

where a_1 is $a_0(\hat{d} = 0.2)$, and $\hat{u}(t_2')$ is limited by (13) which leads to

$$\frac{1}{2}\kappa\mathcal{B}^{max}N(N-1) = \hat{u}(0)\cdot(1-a_0)^{t_1}\cdot(1-a_1)^{t_2'} \tag{23}$$

and solving for t_2',

$$\begin{aligned}t_2' &= \frac{\ln\left(\frac{\sqrt{\pi\kappa}}{2\sigma}\right) - t_1\cdot\ln(1-a_0)}{\ln(1-a_1)}\\ &= \frac{\ln\left(\frac{\sqrt{\pi\kappa}}{2\sigma}\right) - \ln\left(\frac{\sqrt{\pi\alpha\kappa}}{2\sigma}\right)}{\ln(1-a_1)}\\ &= -\frac{\ln(\alpha)}{\ln(1-a_1)}\end{aligned} \tag{24}$$

and finally the second estimator of t^{max} is

$$\begin{aligned}t_2 &= t_1 + t_2'\\ &= \frac{\ln\left(\frac{\sqrt{\pi\alpha\kappa}}{2\sigma}\right)}{ln(1-a_0)} - \frac{\ln(\alpha)}{\ln(1-a_1)}\end{aligned} \tag{25}$$

If we set $\alpha = 1$ then $t_2 = t_1$ as expected. So that α can be used to help scale the predictor for more accurate estimates.

As a final remark on time and effort estimation. Since the effort is a sum over all events where agents are busy over the time t^{max} as per (6), it implies

$$e^{max} \le N\cdot t^{max}. \tag{26}$$

3.7 Maximum Time and Effort

An interesting consequence of (21) is that it places a limit on the time to reach consensus no matter what the group size. That is, for a fixed number of topics, the time to reach consensus has an upper limit. Consider (19) in the limit $N \to \infty$,

$$\lim_{N\to\infty} a_0 = \frac{(2-\sqrt{3})\hat{d}}{3\mathcal{B}^{max}}$$

so that (21) with $\alpha = 1$ gives

$$\lim_{N\to\infty} t_1 = \frac{\ln\left(\frac{\sqrt{\pi\kappa}}{4\sigma}\right)}{\ln\left(1-\frac{(2-\sqrt{3})\hat{d}}{3\mathcal{B}^{max}}\right)} \tag{27}$$

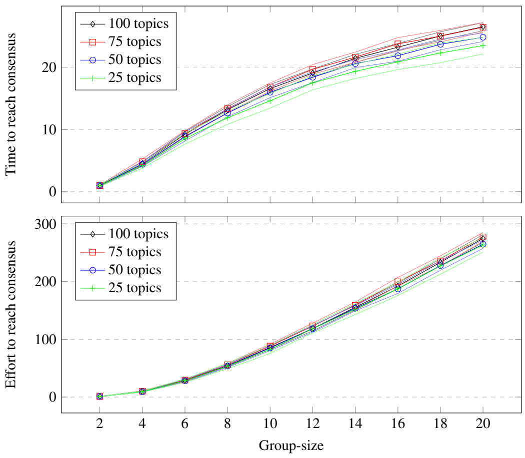

Figure 6: Time (Top) and effort (Bottom) to reach consensus for small groups. Results were scaled relative to a 2-group.

and if $\mathcal{B}^{\max} > 2\hat{d}$ then

$$\lim_{N\to\infty} t_1 \approx \frac{3\mathcal{B}^{\max}}{(2-\sqrt{3})\hat{d}} \cdot \ln\left(\frac{4\sigma}{\sqrt{\pi}\kappa}\right). \qquad (28)$$

The same is not true of the effort, $e^{\max}$, which, from (26), grows at least linear with N and thus is unbounded even if $t^{\max}$ has an upper limit.

$\hat{d}$ has a physical limit in that no matter how fast people talk in meetings, they can only discuss a limited number of topics in a given time. So that $\hat{d}$ is fundamentally constrained. That is not the case with $\mathcal{B}^{\max}$ which represent the complexity of the problem. For large projects, such as the construction of a new factory, or a new industrial complex, $\mathcal{B}^{\max}$ will be very large.

3.8 Assumptions of Cooperation

Agents are fully cooperative, always reach consensus on a topic at a meeting, and topics are independent of each other. This is ideal world assumptions. The results above, in particular the estimation of total effort ($e^{\max}$), total time to reach consensus ($t^{\max}$) and the time to reach consensus independent of number of agents, ($\lim_{N\to\infty} t_1$) are lower limits for real-world consensus processes. That is, no matter how hard real-world agents work, no matter how cooperative they are, no matter how trivial the problem is in terms of interdependencies, a project or task cannot be completed in less time and with less effort than what is given by these results.

In Figure 6 the graphs had been scaled so that the time and effort for a 2-group is 1, and thus all other numbers are relative to this 2-group. Using this as a reference, and considering the results shown in this figure, if a 2-group reaches consensus on a number of topics in a certain time, then, a 6-group will take about ten times longer, and the effort to reach consensus for a 10-group will be about 100 times greater than for a 2-group. In reality, where people may not attend meetings, where they may not reach consensus so easily, and where topics have interdependencies that complicate the process of reaching consensus, it will take even longer. The results presented here should be considered a lower bound on the time and effort to reach consensus due to the simplification assumptions that were made.

4 MODELLING RESULTS

This section presents the results of experiments to identify characteristics of this model through simulations. The following topics are discussed; the experimental setup & data collection, the consensus & entropy measures and their characteristics, the effect on the number of topics, group size, and the presence of artefacts on the time to reach consensus.

4.1 Experimental Setup and Data Collection

The simulator we constructed is primarily for the investigation of team structure, organizational structure, project delivery methodology, and other such organisational aspects. However, this paper only reports on the simulator's design and results. We restrict the organizational structure to polyarchies and investigate the effects of group size, number of topics, the presence of artefacts, and the number of facts that the artefacts contain. The primary measures used are the effort ($e^{\max}$) and the time to reach consensus ($t^{\max}$).

4.2 Consensus and Entropy

Figure 7 (Top) shows the consensus values $u(t)$ over time for twenty simulations of a group of ten agents with $\delta_{ij} = 1,\ \forall i,j \in \{1,2,\ldots,10\}$ and ten topics ($\mathcal{B}^{\max} = 10$). It also shows the average $\bar{u}(t)$ over time averaged over 25000 simulations. Each of these simulations reaches consensus as measured by the discussion in section 2.6, but at different times.

Completing many such simulations allow the computation of a histogram of $t^{\max}$. Figure 7 (Bottom) shows this histogram ($\mu = 48.53$, $\sigma = 5.63$, n=25000) as well as a Normal ($\mathcal{N}$) distribution with the same (μ, σ) parameters. The histogram is not Normal. Visual inspection shows fat-tailed distribu-

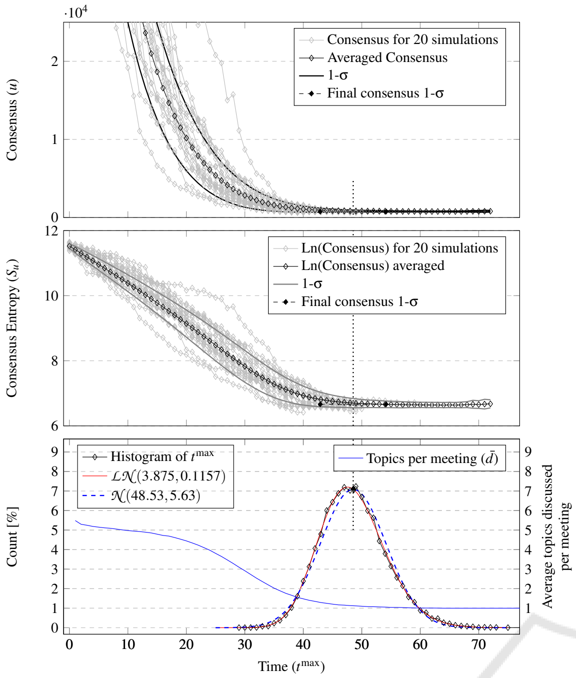

Figure 7: (Top) Various simulations of the 10-group showing the consensus measure over time. (Middle) The same data as in top graph, but now using ln(consensus). (Bottom) Histogram of the time it takes to reach consensus over many such runs ($\mu = 48.53$, $\sigma = 5.63$, n=25000) and Normal and Lognormal fits to the histogram data.

tions and Kolmogorov-Smirnov (KS) and Shapiro-Wilk (SW) tests fail. Figure 12 shows the histogram for a 15-group together with a Gaussian distribution for further edification.

Figure 7 (Middle) shows the same data as (Top) but using the entropy measure. Under this measure the entropy initially follows an approximately linear decrease until very close to consensus. However, the agents then take a significant time to resolve the small differences in views to finally reach overall consensus.

Figure 8 (Top) shows the consensus entropy (S_u) profiles averaged over n=25000 simulations for group sizes of $N \in \{10, 20, \ldots, 50\}$. It also shows (Bottom) the histograms for each of the distributions of final consensus time ($t^{\max}$). As was already hinted at, the distribution is not Normal and the best fit we found was Lognormal $\mathcal{LN}$, even so, it still fails Kolmogorov-Smirnov tests.

4.3 Topics

To characterise the effect of the number of topics per agent, $\mathcal{B}^{\max}$, on the effort and time to reach consensus a number of simulations were conducted for a range of topics (10, 20, ... 100, 200, ... 500) keeping

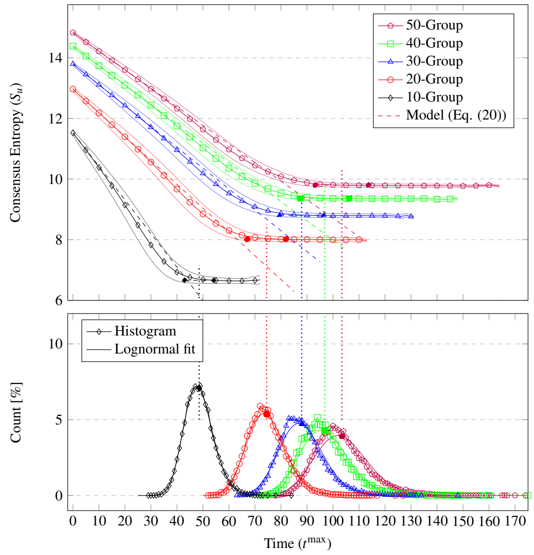

Figure 8: (Top) Group consensus entropy for group sizes of 10, 20, 30, 40, and 50, the bands indicate 1-σ, n=25000. Every 5[th] symbol is shown to avoid symbol clutter. The solid symbols indicate a 1-σ spread in $t^{\max}$. (Bottom) Histograms for the group sizes, and the distribution mean indicated by solid symbols and dotted lines.

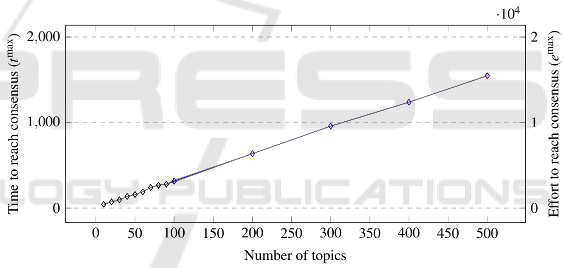

Figure 9: The effect and time to reach consensus as a function of the number of topics.

the number of agents constant (N=10) in a fully connected configuration, n=1 simulations per data point. The results show an expected linear increase in effort and time to reach consensus, see Figure 9 which shows the effort as a function of the number of topics. It would be worthwhile, in future work, to explore the effect of interdependent topics on the time to reach consensus.

4.4 Group-Size

Next we report on the effect of groupsize on the effort and time to reach consensus. We vary the number of agents (N) per group, for group sizes $N \in \{5, 10, \ldots, 100, 200, \ldots, 1000\}$. For each such configuration we compute a number of stochastic simulations to determine mean and standard deviations.

Figure 10 shows effort as as function of group size

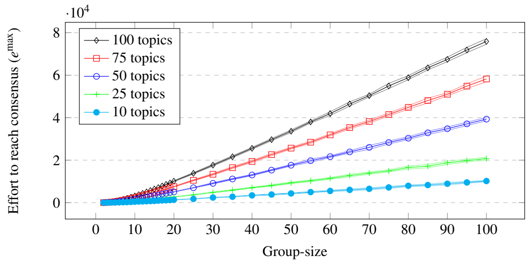

Figure 10: Effort to reach consensus for different group sizes. The bands indicate 1-σ, n=20.

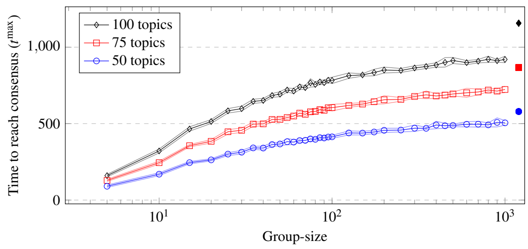

Figure 11: Time to reach consensus for different group sizes. The bands indicate 1-σ, n=10. The solid symbols on the far right indicate maximum consensus time as $N \to \infty$ as given by (27).

and for many different numbers of topics ($\mathcal{B}^{max}$). The *effort* to reach consensus agree with (26) and is linear in N for large $N > 20$.

The *time* to reach consensus as a function of group size has a much more complex relationship as shown in Figure 11. The predictive functions developed earlier, (21) and (25), shows good approximations for t^{max} at least for larger group sizes as is shown in the graph.

4.5 Artefacts

In this subsection we investigate the effects of introducing artefacts that capture views on the range of topics. We are interested in two main questions for the purposes of this paper; firstly, what are the effects of the introduction of an artefact on the effort and time to reach consensus, and secondly, what are the effects of the completeness (in terms of number of topics covered) relative to the number of topics on which agents need to reach consensus.

We compare effort and time to reach consensus for various group sizes by simulating the groups with an artefact and without an artefact. The number of topics are kept the same for agents and artefacts.

Figure 12 shows the entropy profile (Top) averaged over many (n=25000) simulations for a 15-group without an artefact as well as with an artefact. In the case of the group with the artefact, the agents will prioritize working on the artefact above attending meet-

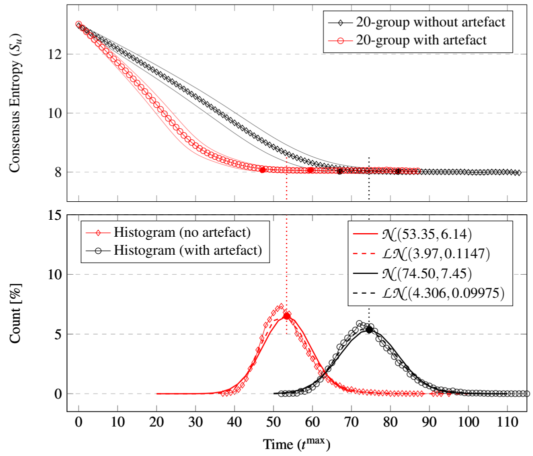

Figure 12: (Top) Consensus Entropy with, and without a supporting artefact. The black band indicate a 1-σ interval. (Bottom) Associated histograms of time taken to reach consensus, without artefact, $\mu = 74.50$, $\sigma = 7.45$, n=25000, and with artefact $\mu = 53.35$, $\sigma = 6.14$, n=25000.

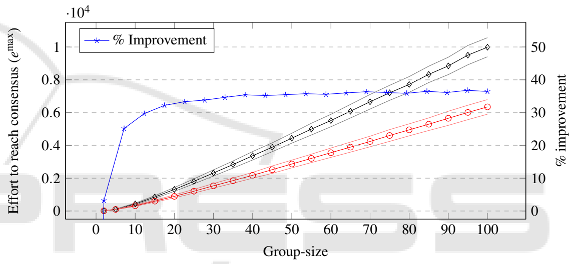

Figure 13: (Top) Time to reach consensus, with and without a supporting artefact. (Bottom) Effort to reach consensus. The bands indicate 1-σ, n=500.

ings.

The same figure (Bottom) shows the histograms for time to reach consensus (with Gaussian fits) which allows the computation of the expected improvement in effort and time. In particular, for this specific example (15-group), the improvement is significant, $t_\Delta = (74.50 - 53.35)/74.50 = 28.4\%$.

This result raises the question whether this improvement can be achieved for all group sizes. We repeat the experiment for many group sizes and find that the addition of an artefact significantly improves the ability of the group to reach consensus (except for very small groups discussed below).

Figure 13 displays the results from repeated experiments with group sizes ranging from 5, 10, ..., 100, both with and without an artefact. Also plotted are the % improvement which shows a consistent (though not constant) improvement of approximately 30% in the time to reach consensus.

An interesting result occurred for small group sizes. The data suggests that groups smaller than five agents reach consensus faster without an artefact. A

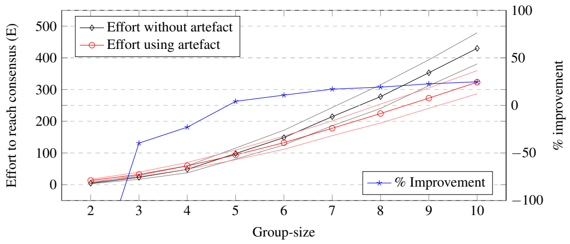

Figure 14: (Effort to reach consensus with and without an artefact. The bands indicate 1-σ, n=500.

4-group needs (on average) 42.8 actions without an artefact and 60.4 actions with an artefact to reach consensus, a 3-group needs 20.9 and 31.3 actions without and with an artefact, and a 2-group needs 3.3 actions without and 12.6 actions with an artefact.

On the other hand, any group with size larger than 5 show significant improvement in time t^{max} when using an artefact, as can be seen from the % improvement plotted in Figure 13 and Figure 14.

5 CONCLUSION

In this paper we described a simulator for studying organisational structure with the aim to model complex organisations and the effects of team structure, organisational structure and the use of artefacts to improve project delivery. We presented theoretical and statistical models for polyarchical structures. We presented the simulation results for modelling polyarchical organisations of various sizes.

Some of the interesting results we found was that for a given problem size, a team of 6 will need approximately 10 times longer to reach consensus what would a team of 2. A team of 10 will need 100 times the effort to reach consensus compared to a team of 2. The use of artefacts to facilitate consensus discussions greatly improve the time and effort needed to reach consensus if the group is bigger than 5. Finally, if the problem has a fixed size, then there is an upper bound on the time needed to reach consensus, no matter how many people are involved (on the assumption that every one is cooperative).

REFERENCES

Birko, S., Dove, E. S., and Özdemir, V. (2015). Evaluation of nine consensus indices in delphi foresight research and their dependency on delphi survey characteristics: a simulation study and debate on delphi design and interpretation. *PloS one*, 10(8):e0135162.

Carter, D. R., DeChurch, L. A., Braun, M. T., and Contractor, N. S. (2015). Social network approaches to leadership: An integrative conceptual review. *Journal of Applied Psychology*, 100(3):597–622.

Chang, M.-H. and Harrington, J. E. (2000). Centralization vs. decentralization in a multi-unit organization: A computational model of a retail chain as a multi-agent adaptive system. *Management Science*, 46(11):1427–1440.

Christensen, M. and Knudsen, T. (2010). Design of decision-making organizations. *Management Science*, 56(1):71–89.

Den Boon, A. K. and Van Meurs, A. (1991). Measuring opinion distributions: An instrument for the measurement of perceived opinion distributions. *Quality and Quantity*, 25(4):359–379.

Jones, S. L. and Shah, P. P. (2016). Diagnosing the locus of trust: A temporal perspective for trustor, trustee and dyadic influences on perceived trustworthiness. *Journal of Applied Psychology*, 101:392–414.

Keupp, M. M., Palmié, M., and Gassmann, O. (2012). The strategic management of innovation: A systematic review and paths for future research. *International journal of management reviews*, 14(4):367–390.

Kian, M. E., Sun, M., and Bosché, F. (2016). A consistency-checking consensus-building method to assess complexity of energy megaprojects. *Procedia-social and behavioral sciences*, 226:43–50.

Lang, J. W., Bliese, P. D., and de Voogt, A. (2018). Modeling consensus emergence in groups using longitudinal multilevel methods. *Personnel Psychology*, 71(2):255–281.

Reagans, R., Miron-Spektor, E., and Argote, L. (2016). Knowledge utilization, coordination, and team performance. *Organization Science*, 27(5):1108–1124.

Roselló, L., Prats, F., Agell, N., and Sánchez, M. (2010). Measuring consensus in group decisions by means of qualitative reasoning. *International Journal of Approximate Reasoning*, 51(4):441–452.

Sáenz-Royo, C. and Lozano-Rojo, A. (2023). Authoritarianism versus participation in innovation decisions. *Technovation*, 124:102741.

Wei, Q., Wang, X., Zhong, X., and Wu, N. (2021). Consensus control of leader-following multi-agent systems in directed topology with heterogeneous disturbances. *IEEE/CAA Journal of Automatica Sinica*, 8(2):423–431.

Will, M. G., Al-Kfairy, M., and Mellor, R. B. (2019). How organizational structure transforms risky innovations into performance–a computer simulation. *Simulation Modelling Practice and Theory*, 94:264–285.

Yan, H.-B., Ma, T., and Huynh, V.-N. (2017). On qualitative multi-attribute group decision making and its consensus measure: A probability based perspective. *Omega*, 70:94–117.

Young-Hyman, T. (2017). Cooperating without co-laboring: How formal organizational power moderates cross-functional interaction in project teams. *Administrative Science Quarterly*, 62(1):179–214.