EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling

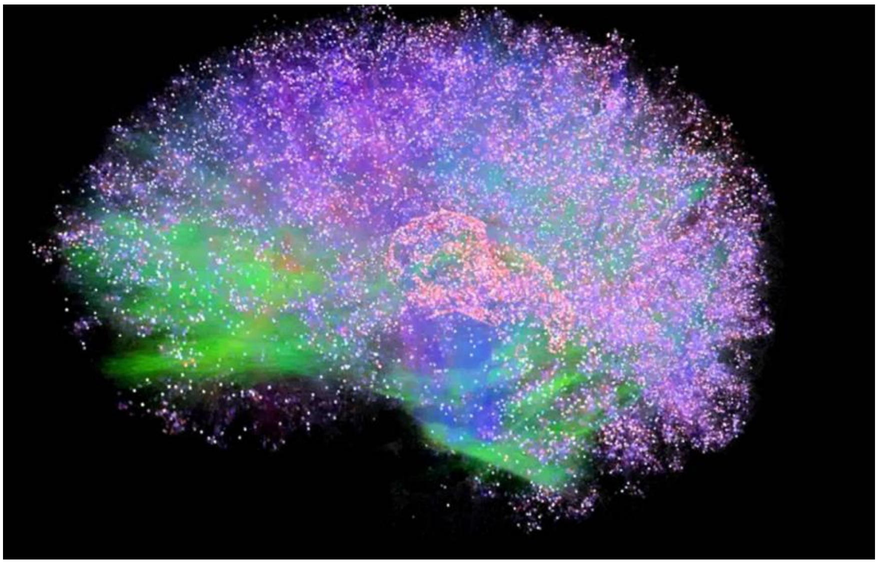

Integrated Information Theory for Organisational Consciousness

Jonathan Peters

Contents

1. Introduction

Effective communication and information flow are crucial for the smooth running of an organisation. If members of an organisation are overwhelmed with information, it may prevent work being done, yet the absence of communication is clearly an obstacle that prevents members from doing their work. Often, a simple top down communication model has been used. Workers communicate with their managers, who in turn ensure that their workers are complementing each other. We show an example of this type of communication model in Figure 1. However, there are flaws in this strictly hierarchical approach.

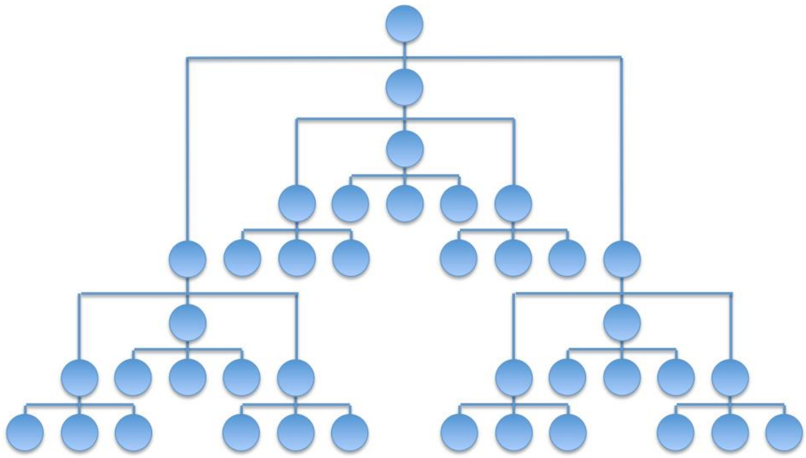

Figure 1: An example of a completely hierarchical business communication model that features a boss at the top, who interacts with managers. These managers then interact with sub managers, who in turn interact with the workers. Hence for a message to get from one worker to another, it must travel through multiple layers of the hierarchy. No horizontal communication occurs in this case.

This top down approach assumes that the work of the company can be subdivided into distinct departments which don't interact with each other, which is usually never the case. Even if different departments focus on completely different problems there will usually be an overlap in what they do. For example, as telecommunication technology evolves, BT needs to shut down parts of the network which are no longer needed. This is a complicated problem, and spans the interest of many of BT's departments. If these departments act independently from each other as they shut down parts of the network, the action of one might negatively affect the others. It is clear that this situation would be better if the departments whose actions might affect each other would coordinate in order to determine the best course of action.

We approach the problem of understanding communication channels in a company using Integrated Information Theory, a novel approach to computing consciousness.

A key issue is to determine the optimal level of cooperation and communication between the departments. In particular we need to address how to determine which departments need to communicate, how frequently they should do so, and whether communication is redundant in specific scenarios so that departments can operate independently of each other. Our aim is to evaluate Integrated Information Theory (IIT) to see if it can address some of these challenges.

IIT is a theory developed for the purpose of measuring consciousness within the brain, which is a biological example of large scale, complex yet efficient, communication. Our goal is to model communication flow within an organisation as a network, and test whether IIT can capture different quality metrics. In practice, this will involve thinking about the entire organisation as if it were the brain and individuals within the organisation (or teams of individuals) as the neurons.

2. The Task Allocation Problem

Within the broad context of communication within an organisation, we focus on the specific example of task allocation. We assume that the organisation has a number of teams, each with different capabilities, and we aim to determine how they should decide who does what task, and how should this be accomplished when not all teams are in constant communication with each other.

Assumptions of Model

We make the following assumptions:

- The organization generates revenue by completing tasks and wants to complete as many tasks as possible.
- There are n different types of skills. Each team possesses one or more of these skills, and each task requires one or more in order to be completed.
- Each team can only be working on a single task at any given time. A task may have multiple teams working on it.
- In isolation, a team only knows what skills it has, whether there is a task to be completed, and what skills that task requires. It does not know that any other teams exist, or what skills they have. If a team becomes 'connected' to another, then it knows what that neighbouring team's skills are, and whether it is currently busy or not.
- Teams decide whether they should work on a task by choosing the optimal arrangement, based on the amount of information that is available to them.

We consider a company that seeks to maximize profits by assigning tasks to its workers efficiently.

We give an example of a communication graph in Figure 2. Each node refers to a team. If there is an arrow pointing from team A to team E, this teams that team A is listening to team E, and hence knows what team E's abilities are and whether team E is working. However, team E is not giving attention back to team A. Only if there is a two way arrow, like between teams A and C does two way communication exist.

When two teams are connected to each other, then that means that they are communicating.

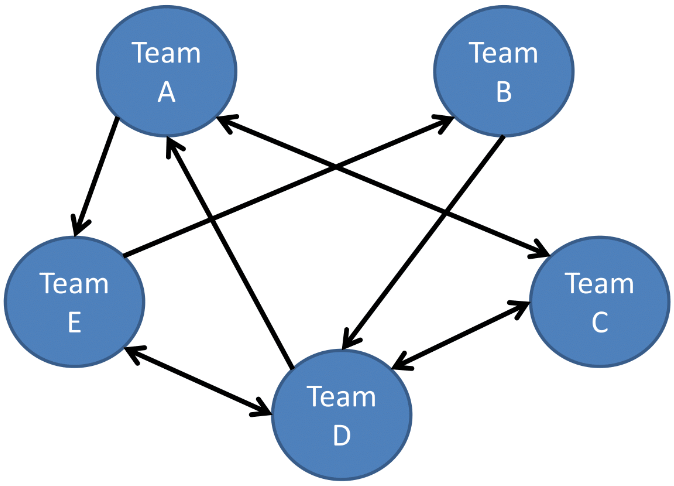

Figure 2: Schematic showing an example of a communication graph between 5 teams in an organization. The arrows indicate which teams are listening to each other.

We quantify the effectiveness of the communication structure of the organization by measuring how often things go wrong. We identify 3 particular scenarios to avoid:

1. Cost of Opportunity: A task is ignored even though there are teams available which could do the task, because the teams think that someone else will do the task.

2. Redundancy: Two or more teams do the same task.
3. Wasted Skills: An overqualified team is assigned to a task.

As scenario 1 and 2 are the wastes that result directly from miscommunication, we view them as the most relevant to measure. We consider them to be far worse than scenario 3.

Team Assignment Process

The next step is to determine how teams are assigned to incoming projects. In order to avoid the three wasteful scenarios, teams need to check the following:

A team will work on a task, if it considers that is the most efficient way for task to be completed given the information available to the team.

1. Can they do the task? If so, ensure the task is completed.
2. If not, can they see another idle team with whom they can collaborate to the project?
3. Is there a less qualified neighbouring team that could do the project instead? If so, minimize waste and avoid repetition by letting the less qualified team do the project.

A team can only apply these policies based on the information it has, and this depends on which teams are communicating with each other. With any communication structure, short of everyone talking to everyone, there will be waste.

We show several communication examples in Figure 3. We consider the teams A, B, and C, as shown in Figure 3 (left), where A is more skilled than B who is more skilled than C. If a task arrives which all three of them can do, then A will not do the task, because it can see that B is more suited. Meanwhile B also won't complete the task, because C is more suited. Finally C will actually do the task. Hence with 2 edges between the teams, the situation is handled correctly. However if we consider the teams D and E as shown in Figure 3 (right), if D is more skilled than E and a task arrives which both E and D can do, this will result in both of them doing the task. This is because D doesn't know anyone else will be doing it, while E sees that, being the less skilled of the two, it should take the task.

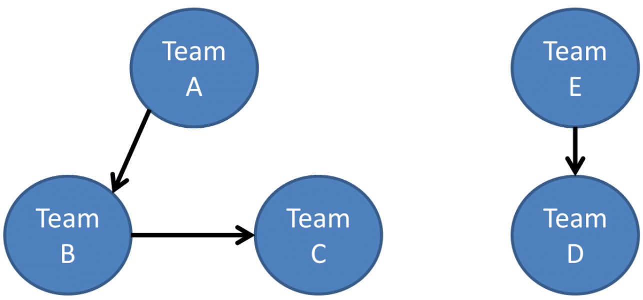

Figure 3: Schematic showing two possible communication examples. The arrows indicate the direction of communication.

Integrated Information Theory

The *integrated information* of a network is the information it specifies above and beyond the sum of its parts.

We use Integrated Information Theory to see what insights it may offer about the operation and structure of organizations. What is IIT? It defines a concept called integrated information, which measures how much information a system of nodes, e.g. neurons, has above and beyond the sum of its individual parts. The developers of this theory argue, mainly using philosophical reasoning, that integrated information is the same thing as consciousness.

We apply IIT to this problem by modelling an organization as we would model the brain: as a network where a dynamical process occurs, thinking of the individual teams as we would neurons. Then we compute the integrated information of the organization. Our aim is to determine whether the integrated information relates to the cost of waste.

One motivation for considering integrated information is the way it requires feedback. If communication is all unidirectional, then the organization will have zero integrated information. Furthermore, IIT distinguishes important information from redundant information. This is a promising feature that we might want to exploit, as we want to ensure that only teams that need to communicate do so.

3. Results

We assign numerical values to the waste incurred by each possible scenario. Then by computing how often we expect each scenario to occur, we determine the cost of waste. This is the first measure which we use to quantify the effectiveness of the organisation. The other measure we use is the Integrated Information generated by the network.

We investigate the relationship between these two values by changing who talks to whom, i.e. the structure of the organisation. We focus on an initial five team scenario, where the first 3 teams are all talking to each other, and the last two are talking to each other. Then we allow 4 additional connections to be made. We consider all possible ways that this can be accomplished, of which there are 495, and plot the results in figure 4. By doing so we investigate how the two clusters of workers should share information, when it is expensive to do so, in the most efficient way. We see in figure 4 that there is a correlation between integrated information and the cost due to miscommunication. In general, we see that higher values of integrated information result in less waste.

Maximizing Integrated Information is related to minimizing the waste due to poor communication.

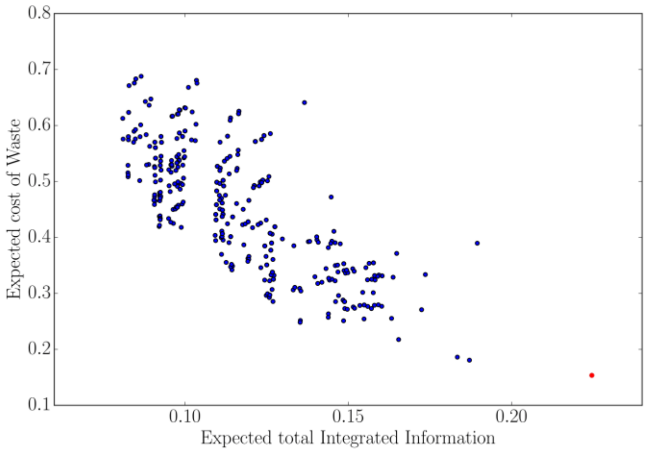

Figure 4: Graph showing the expected cost of waste and the expected total Integrated Information for each choice of 4 new communication links.

In figure 5, we present the structural design which both minimizes waste, and maximises Integrated Information. The original links are coloured in black, and the 4 new links are in red. This graph corresponds to the red data point in figure 4. In particular, we observe that the smaller cluster ED must pay greater attention to the larger cluster ABC than vice versa. This occurs with D in the small cluster listening to every node in the large cluster.

In general, two way conversations reduce the costs of missed opportunities. However, for avoiding unnecessary duplication of labour, unidirectional communication is more efficient. Therefore, the optimal structure in general will depend on what an organisation views to be worse: missing an opportunity, or needless repetition of labour.

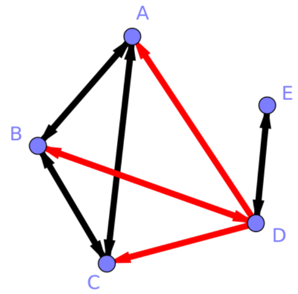

When adding four communication links, we see a combination of unidirectional communication and two way conversations resulting in the optimum.

Figure 5: The structural design which optimized costs. The red lines are the new communication channels.

Consolidation of Integrated Information Theory

Since first being proposed in 2004, many papers have been published on Integrated Information Theory. In each paper and new version of IIT (the most recent version is 3.0), the notation and conventions are changed. In addition, even though the calculation of integrated information is a mathematical operation, in its current state IIT lacks robust mathematical underpinnings. Many of the formulas and procedures outlined in IIT are not rigorously expressed. We investigate the different notations used, determine the most consistent version, and make the adjustments necessary for mathematical consistency.

Having done this, we find that IIT has three main drawbacks:

The expense of computing integrated information makes it an inadequate solution for a moderately sized organization.

- It is a challenge to compute. In practice, the computation of each individual data point in figure 4 took between 30 minutes and 2 hours. Hence, the expense of calculating IIT as currently defined makes it unsuitable to use at the present.
- IIT requires that all possible scenarios that could occur within the organization are understood. In practice, we only know about the scenarios which we have actually observed, and cannot anticipate all possibilities. This restriction makes it infeasible for IIT as currently defined to be applied to a more realistic and complicated model.
- Integrated Information depends both on the structure of the network, and the networks current state. Small changes in the current state may result in large changes in integrated information. Hence, unless we have perfect knowledge of the state which the organization is in, our measurement of integrated information is meaningless. It would be preferable if the measurement depended only on structure.

However, integrated information does capture certain phenomena that are of interest. It ignores redundant information and captures the significance of feedback. The main obstacles are steps which can easily be explained conceptually, but are expensive mathematically. If a method could be found to approximate or skip entirely the computationally expensive steps in IIT, then there may still be use in the calculation.

Expanding the model

In this model, we have studied how poor communication affects the assignment of tasks. However, it fails to capture how the work done on one task might impact the progress on a related task. This feature must be included in order to address BT's original problem: what communication needs to occur to prevent departments interfering with each other. In order to reach this goal, we propose the following extensions to the model:

We should extend our model to allow progress on different tasks to influence each other, as happens in reality

- Record not merely that a task has been assigned, but who is working on it, and the progress made.
- Allow connections between tasks well as teams, showing how work done on one task might affect progress on another.
- Let the rate at which a task is completed depend on the communication between the teams working on the task, and the other tasks which affect it.

- Assign a cost to communication in order to identify communication links which, though beneficial, are outweighed by their cost.

These extra features would not affect the method of computing waste or integrated information; only make the model more realistic. However, the computational time required for this at present makes this may make the approach infeasible.

4. Discussion, Conclusions & Recommendations

We have derived a model for communication flow which captures to effects of poor communication on an organisation. Using this model, we compared the effect of integrated information on waste, concluding that, in general, the more integrated information that a communication arrangement has, the less wasteful it will be. Due to computational effort required and time constraints, we were not able to determine the best communication flow arrangement out of all possibilities, but only out of a select few, even when we only have 5 teams.

Therefore if IIT is investigated by BT in the future, we suggest that all the steps involved should be scrutinized mathematically and remove any steps which may be unnecessary. Furthermore, not all features of Integrated Information Theory are necessarily critical to the communication flow problem (even if they are critical for computing consciousness). Identifying and removing such features would make this approach more tractable.

While the types of waste identified are important, our ultimate aim is to incorporate other types of waste in order to identify scenarios when two departments are actively working against each other without realising it. To this purpose, we suggest expanding the model as described in section 3. In particular we suggest:

- Assigning a cost to communication links;
- Allowing the progress made on projects to affect each other.

Our main conclusion is that, in its current formulation, IIT is too computationally complex to be of practical use in large, complex, organisations. However, using conventional network approaches instead of IIT may yet yield further insight.

5. Potential Impact

Despite our main conclusion that IIT in its current form is not computationally tractable, we have laid a solid mathematical foundation for future investigation.

Steve Cassidy, Chief Researcher – Systems Science, commented "*The creation of data-driven models enable BT to cohere complex and strategically significant decisions. These decisions affect the many diverse areas of the business: operations, finance, strategy HR etc. Finding globally optimal strategies and co-ordinated delivery is not possible without the "organisational intelligence" created by these tools. Along with the tool development, we would like to understand, characterise and optimise the connectivity that such tools offer, so as to maximise the benefits. This study of Integrated Information Theory has proved invaluable in understanding the potential and limitations of this approach. It has rigorously and succinctly identified which elements of the approach should be developed, with new possible formalisms, and which directions look intractable. This has mapped out a set of promising future directions for BT Research, and identified dead ends to be avoided. The first of these add value, and the second avoids waste in future work in this area. An enhanced ability for organisations to design their connectivity for maximum effectiveness, as opposed to relying on the traditional pyramids, would be a great asset to industries and enterprises of all kinds.*"