

SUSTAINED ADAPTABILITY: THE TRANSACTION LEVEL

Jan Maarten Schraagen^{1, 2}

¹ TNO, Soesterberg, The Netherlands

¹ jan_maarten.schraagen@tno.nl

² University of Twente, Enschede, The Netherlands

² j.m.c.schraagen@utwente.nl

Abstract

A new system level, called the 'Transaction Level' is introduced. I argue that such architectures should not be couched in (macro)cognitive terms, but rather in terms of networks of nodes and links that effectuate transactions. The principle of relationality governing this level states that links are selected to attain transactions. The transaction level is a true systems level rather than a perspective on a particular unit of analysis (individual, team, organization). The novelties and advantages of the introduction of the transaction level for the field of resilience engineering are: (1) an increased emphasis on longitudinal data collection and use of social network analysis as one of the tools to analyse data collected on nodes and links; (2) providing an explanation for when transactions fail and may lead to accidents in sociotechnical systems; (3) a renewed emphasis on the study of patterned interactions of sociomaterial assemblages; (4) providing a language for describing architectures for sustained adaptability and thus advancing relative invariants in the study of Layered Networks.

1 INTRODUCTION

Advances in technology and automation have enabled high levels of connectivity, not only amongst individuals, but also amongst teams, organizations and, at global levels, even countries. Networks are a prime example of such highly interconnected units and have been studied as such in a field collectively denoted as 'network science'. From a slightly different angle, a network perspective has been taken in many fields as well, ranging from sociology and psychology to logistics and biology. Recently, in the field of resilience engineering, Woods (2015) has introduced the concept of Layered Networks, referring to the multi-layered character of networks, as well as its numerous interconnections crossing levels.

Layered Networks show 'graceful extensibility' in successful cases of sustained adaptability, that is, these networks have the ability to extend their capacity to adapt when surprise events challenge their boundaries. A key concept in this view is 'adaptation', which includes adjusting behaviour and changing priorities in the pursuit of goals. As surprise occurs continuously, units within the network constantly need to monitor their environment and regulate their resources, in joint coordination with other interdependent units in the network.

Undoubtedly, for both human and artificial systems, knowledge is one of the most important resources adaptive units can bring to bear. Any system that fluently applies knowledge in the service of goals operates at what Newell (1982) has termed the 'knowledge level'. At the 'knowledge level', the principle of rationality applies: if an agent (e.g., an expert) has a goal and knows that knowledge A will bring him or her closer to that goal, then the agent will choose knowledge A. As knowledge is always finite (principle of bounded rationality; Simon, 1955), adaptation can only be local and perspectives of any unit in the network are bounded. In Newell's (1982) terms, the 'knowledge level' is a radical approximation, that is, entire ranges of behaviour may not be describable at the knowledge level, but only in terms of systems at a lower level, i.e. the cognitive level.

However, if knowledge is finite, the only way for Layered Networks to show graceful extensibility is by aligning and coordinating across multiple interdependent units in a network, by shifting and contrasting over multiple perspectives, hence by extending the range of adaptive behaviour of other units. I will argue that this, in

effect, calls for a different system level, right above the knowledge level, which I will call the 'transaction level'.

2 SYSTEM LEVELS

There are many definitions of what a 'system level' is. According to Newell (1982), a true system level is a reflection of the nature of the physical world, not simply a level of abstraction. It should be a specialization of the class of systems capable of being described at the next (higher) level. Aggregation occurs within each level and does not take us to the next level on its own; meaning should be added and some things therefore become invisible at the next level (this is another way of saying that phenomena at higher levels have emergent properties). Therefore, although the levels are ontologically irreducible to each other, each level may still be implemented at the next lower level.

Note that this view of system levels is quite different from what we normally take to be 'units of analysis'. For instance, the distinction between micro, meso and macro levels (individual, group, organization) is not a true system level description, in Newell's definition. For instance, Karsh, Waterson, and Holden (2014) proposed 'mesoergonomics' as a way to specify macro- and microergonomic integration. Their aim was to reveal cross-level interactions and to describe relationships between and among levels rather than describing phenomena that emerge from their components but that cannot be explained by them (Hackman, 2003).

Thus, a system level description should be distinguished from a unit of analysis description. For instance, if we take a particular unit of analysis as our focus, let's say an individual, then, from a system level perspective, we could still describe the individual at multiple system levels simultaneously. The individual could be described at a neurological level, at a cognitive level, at a knowledge level, and so forth upward as well as downward (towards the basic particle level). Basically, this conforms to the division of work in science: physics and chemistry deal with lower system levels than psychology and sociology.

It is the goal of science to find invariants in the behaviour it tries to describe. In physics and chemistry, scientists aim for absolute invariants that hold over all stretches of time. When studying artificial (i.e., goal-directed, hence adaptive) systems, such as Layered Networks or in general human-technology-organisation ensembles, these absolute invariants do not hold up. As noted by Simon (1980), adaptive or artificial systems will change as their environment changes. Therefore, it is difficult to discover and verify empirical invariants in these systems, as any laws that govern their behavior must contain reference to their relativity to environmental features. As noted by Simon (1980, p.36): « It is common experience in experimental psychology, for example, to discover that we are studying sociology—the effects of the past histories of our subjects—when we think we are studying physiology—the effects of properties of the human nervous system ».

Simon goes on to state, however, that we should not despair of finding invariants. Rather than looking for absolute invariants, as the physicists do, cognitive science should search for relative invariants that hold over considerable stretches of time and ranges of systems : « What is invariant in adaptive systems will depend on the time intervals during which we observe them » (Simon, 1980, p. 36). Simon then mentions three time scales on which adaptation takes place : the shortest time scale in which heuristic search takes place ; a longer time scale in which learning (storing and retrieving knowledge) takes place ; and the longest time scale during which social and biological evolution takes place, in the form of discovery of new knowledge and strategies and the transmission of one system to another.

Newell (1982, 1990) has later expanded on this view by taking the time scales Simon (1980) described as his starting point. Newell (1990) distinguished a biological, cognitive, rational, and social 'stratum', each defined by a particular time band during which processes take place. For instance, the typical cognitive processes take place between 100 msec and 10 sec and the typical rational processes between minutes and a few hours. However, Newell (1982) went one step beyond Simon and equated these time scales with system levels that occur in nature (in his 1982 article, Newell confined these system levels to computer system levels, but it became apparent in his 1990 book that he meant these system levels to apply to all biological and artificial systems).

Therefore, system levels should be defined according to the time scales they describe, not to the unit of analysis they describe. The field of resilience engineering studies adaptation of Layered Networks over long time scales (days to years). The relative invariants that resilience engineering hopes to describe are architectures for sustained adaptability. The question then arises what these architectures look like and how sustained adaptability should be described.

3 TRANSACTION LEVEL

Most researchers in cognitive psychology are familiar, albeit implicitly, with the cognitive stratum. Memory search processes take place at this level, as well as some higher-level, but still basic, problem-solving processes, such as building up mental representations and carrying out mental computations. For instance, the description of real-time pilot behaviour when dealing with information presented on displays usually takes the form of a cognitive level phenomenon.

Fewer researchers are familiar with what Newell (1982) termed the 'knowledge level'. This is behaviour at longer time scales. At the 'knowledge level', the principle of rationality applies: if an agent (e.g., an expert) has a goal and knows that knowledge A will bring him or her closer to that goal, then the agent will choose knowledge A. From the outside, the behavior of the expert is highly predictable, once we know the expert's goals and the knowledge that is required to attain those goals. Many studies in resilience engineering in fact take place at this level. Descriptions of how individuals and teams monitor, respond, anticipate and learn usually abstract away from computational limitations and lack of knowledge and rather assume that knowledge is readily available in the service of goals. For instance, in our studies on enhancing resilience in railway teams, Willy Siegel and I have focused on collaboration displays that foster reflection by making implicit knowledge explicit. One such collaboration display is the Resiliencer (Siegel & Schraagen, 2017a). The Resiliencer provides weak resilience signals on performance, workload and safety boundaries of the railway system as well as analysis functions to connect system level signals to personal identifiable details (e.g., a specific train). The application allowed a team of rail signallers to analyse movements toward system boundaries and share knowledge on these movements (Siegel & Schraagen, 2017b). Over the course of our observations, rail signallers increasingly analysed and reasoned about their work. This enriched knowledge beyond procedures, enhancing the ability to cope with the unexpected and unforeseen (Siegel & Schraagen, 2017a). The time scales of our observations was indeed from minutes to hours (after-shift reflections lasted approximately 30 minutes). Also, we looked at the enrichment and exchange of knowledge among rail signallers, which again is a purely knowledge-level phenomenon.

Yet, knowledge-level descriptions have their limitations when we aim at describing architectures for sustained adaptability. Although we analysed the rail signaller's team reflection over the course of a one-week period, we don't know whether their behaviour leads to sustained adaptability, that is, resilient behaviour over longer periods of time. As noted by Siegel in his dissertation (2017, p. 133): "The staff level needs to incorporate the knowledge into procedure updates and take structural actions for the medium and long term, with all the rail parties. This implies an organisational structure change which will influence the evolution of needed skills in the whole organisation and eventually will result in resilience enhancement of the rail sociotechnical system." Yet, the issue is not just one of a difference between the team level and the staff level. As mentioned above, a team could also be observed over a period of months, and a staff could be observed over a period of hours. A unit of analysis (team/staff) is not the defining difference between system levels. What is the defining difference is that meaning should be added to a particular level, which results in some things becoming invisible and other things becoming visible (because of different concepts being used). Hence, as long as we speak of knowledge being incorporated into procedures at the staff level, we are still describing phenomena at the knowledge level. To move to a different system level, we need a different vocabulary.

I now propose that there does exist yet another system level, which I will call the transaction level. It is a true systems level in Newell's sense, that is, it is a reflection of the nature of the physical (and social!) world and not just a point of view that exists solely in the eye of the beholder. It is not a level of analysis that can be applied to any unit of analysis, as, for instance, a network perspective that can be applied to both brains and societies. It is not an aggregation of knowledge, so its behavior cannot be obtained by expanding agents at the knowledge level with simply more knowledge. Rather, no amount of knowledge added will yield the transaction level properties that are characteristic of this level.

A quick overview of the transaction level, in terms of the system under consideration, its components, its laws of composition, its behavior laws and its medium, are in order before entering into details.

The system at the transaction level, the entity to be described, is the network. The system's primitive elements, its components, are nodes and links. Thus, a network is composed of a set of agents and a set of links. The components are assembled into systems by laws of composition that yield strength and reciprocity. The medium at the transaction level is the transaction (as might be suspected). Thus, the network generates transactions by connecting nodes through links. The transactional content may differ widely, from affect and influence to goods and services, and information. Finally, the behavior law, how the system depends upon its components and composition, is the principle of relationality: links are selected to attain transactions. As links

are characterized by strength and reciprocity, the generation of transactions is dependent upon link strength and reciprocity.

In contrast to the knowledge level, the concept of ‘goal’ does not play a role at the transaction level. However, just as with the knowledge level, the transaction level is a radical approximation : entire ranges of behavior may not be describable at the transaction level, but only in terms of systems at a lower level. For instance, the transaction level is poor for predicting how team members that have never met before will interact. It is also poor for predicting the effectiveness of the introduction of a new technology on a person’s behavior. However, it is good for predicting the impact of losing someone central to an organization’s informal network. It is also good for predicting that a well-established team will exchange relevant information in a timely fashion.

The physical structure of a transaction is filled indirectly and approximately by knowledge systems at the next lower level. This is depicted in figure 1 :

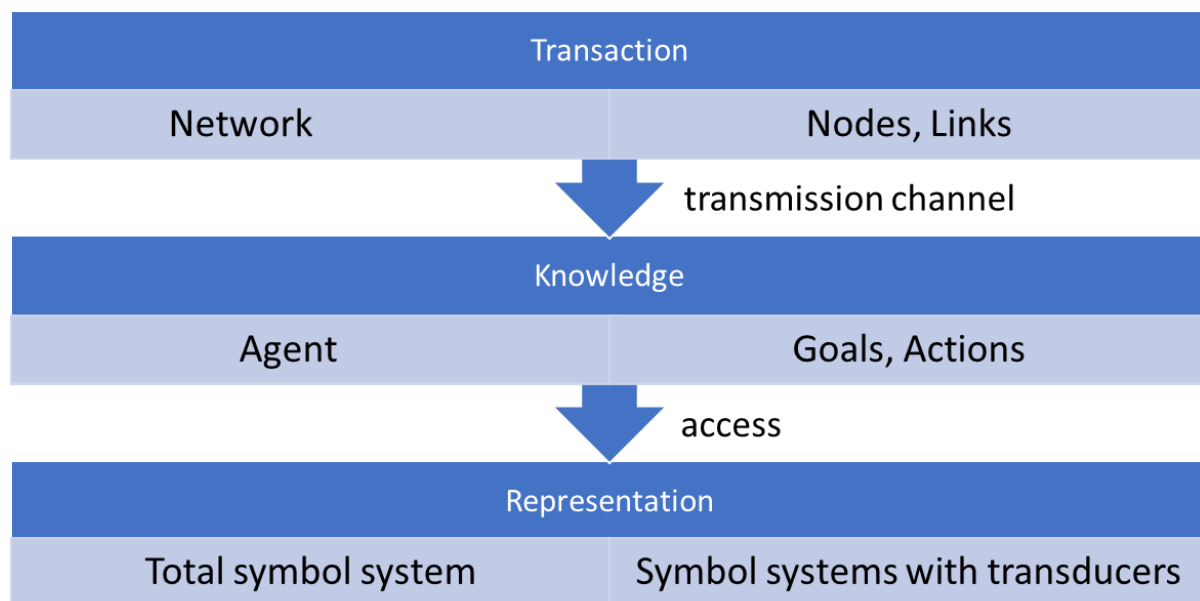


Figure 1. Reduction of the transaction level to the knowledge level and the symbol level

The slogan equation for connecting the transaction level to the knowledge level is: Knowledge = Transaction + Transmission channel. What this means in psychological terms is that humans are social beings that have knowledge of all kinds of possible transactions, but only to a certain extent, due to limitations on communication bandwidth. This knowledge is a person’s social capital, or the interpersonal relationships they have with others that enable successful functioning (Hollenbeck & Jamieson, 2015). Social capital is conceptually distinct from the individual-level constructs and it has the potential to enhance prediction of individual and team performance incrementally over those constructs. Social capital information can be quantified using social network analysis (see Schraagen, 2015, for a discussion of the relation between resilience and social network analysis).

4 SIMILARITIES AND DIFFERENCES WITH OTHER APPROACHES

The emphasis on the exchange of transactions between actors in a network is very similar to what Stanton et al. (2006) described in their theory of distributed situation awareness. These authors took a systems level approach and noted that what mattered was that the right information was passed to the right agent at the right time, not that all information is available with a single human agent. Neville et al. (2016) also used the concept of ‘transaction’ in a similar, though more restricted, way than I do. They referred to transactions as an exchange of situation awareness between agents, which is more than the mere communication between agents. Transactions are enriched by specific and individual interpretations of each agent and so may provide a clue to other agents as to what one individual is working on. As transactions hold the key to safe and efficient performance, post accidents this means that investigators need to understand not only what information was lost but also what transactions were inadequate or were required but not forthcoming (Salmon, Walker, & Stanton, 2016).

However, although superficially similar, Stanton and co-workers' approach differs from my description of the transaction level in several ways. First, although Stanton et al. correctly describe situation awareness as a system level phenomenon rather than an individual level phenomenon, a 'system level' in Newell's sense is not the same as a 'system level' in Stanton's sense. What Stanton et al. are referring to is a difference in the unit of analysis, namely the individual versus the 'system' (where 'system' refers to a sociotechnical system). Hence, in their view, it makes sense to say that a 'system' holds situation awareness or loses situation awareness, whereas in my view, at the transaction level, the concept of situation awareness simply does not exist, as it exists at the knowledge level.

Second, Salmon, Walker, and Stanton (2016), in their application of their Distributed Situation Awareness theory to the Air France 447 accident, focused on the sharp end of the communication patterns of the pilots involved. Again, although valuable, this is a knowledge level analysis rather than a transaction level analysis. What needs to be determined, if we are looking for architectures for sustained adaptability, is what underlying, structural relational patterns were present in the cockpit that led to this specific exchange of information at this particular moment in time. Furthermore, these authors restrict transactions to informational transactions, whereas power relationships and affective relationships also play an important role. Thus, rather than focusing on who said what to whom at a particular moment in time, we should focus on structural affective, power and informational relationships between human and non-human agents in the network.

I agree with Salmon, Walker and Stanton (2016) that failed transactions lie at the root of accidents occurring in complex sociotechnical systems. However, this statement does not explain how transactions fail. The distinction between the transaction level and the knowledge level does at least begin to explain how transactions fail: first, by structural differences in link strength and reciprocity; second, by restrictions on communication bandwidth. As to the first aspect, one needs to ask whether team members have been able to develop links with each other, and what power relations they have with each other. Of course, as all system levels simultaneously play a role at all times, and the transaction level, just as the knowledge level, is a radical approximation, this implies that, for example, when team members do not know each other well and have not built up structural links with each other, the transaction level is poor for predicting how these team members will interact. We need to invoke the knowledge level in that case to explain absent, inappropriate, incomplete or misunderstood transactions. Invocation of the knowledge level is related to the second aspect, restrictions on communication bandwidth. Even if individual team members possess all relevant knowledge, they may be unable to share all that knowledge, particularly in stressful situations, due to the fact that they can only speak, see and hear so much. A team of experts is not by definition an expert team, if team members don't know when to communicate what information to whom, or don't dare to communicate. This goes back to the first aspect of the establishment of structural and reciprocal links between team members (or agents in general: the interactions between humans and non-humans may also be fruitfully analysed in terms of link strength and reciprocity).

A second field that should be noted is the field of organization studies and the sociology of technology. In particular, Actor-network theory (Latour, 2005) and the concept of 'sociomateriality' (Orlikowski & Scott, 2008) have been highly influential in eliminating the conceptual distinction between humans/organizations and technology. Their view of the social and technical worlds is described by Orlikowski and Scott (2008, . 457) as: "Humans/organizations and technology are assumed to exist only through their temporally emergent constitutive entanglement." As in the transaction level, the principle involved is the principle of relationality and the system to be described consists of 'sociomaterial assemblages'. People and things only exist in relation to each other; in fact, Latour (2005) goes so far as to claim that the distinction does not exist in the empirical world, but rather is an invention by scholars to demarcate disciplines of study. Sociomateriality may be defined then as "[e]nactment of a particular set of activities that meld materiality with institutions, norms, discourses, and all other phenomena we typically define as "social."" (Leonardi, 2012, p.42).

For resilience engineers, the novelty of these approaches lies not so much in the recognition of the importance of technology in the social world, nor even in the recognition that materiality is present in each and every phenomenon that resilience engineers consider 'social'. Rather, the novelty may be in the absolute elimination of any distinction between humans and technology. This is in complete alignment with the description of the primitive elements of the transaction level. The distinction between nodes and links does not make a difference between human nodes and technological nodes. In fact, with the increase in intelligence in technological artifacts, the distinction between humans and technology has already faded. This is not to say that humans and technological artifacts are ontologically identical and indistinguishable. Rather, at lower levels they become distinct entities (perhaps already at the knowledge level, but certainly at the cognitive

level). However, at the transaction level, the distinctions are irrelevant.

5 CONCLUSIONS AND IMPLICATIONS

I have advanced the notion of the transaction level as a separate system level directly above the knowledge level. The novelties and advantages of the introduction of the transaction level for the field of resilience engineering are: (1) an increased emphasis on longitudinal data collection and use of social network analysis as one of the tools to analyse data collected on nodes and links; (2) providing an explanation for when transactions fail and may lead to accidents in sociotechnical systems; (3) a renewed emphasis on the study of patterned interactions of sociomaterial assemblages; (4) providing a language for describing architectures for sustained adaptability and thus advancing relative invariants in the study of Layered Networks.

REFERENCES

- Hackman, J.R. (2003). Learning more by crossing levels: evidence from airplanes, hospitals, and orchestras. *Journal of Organizational Behavior*, 24, 905-922.
- Hollenbeck, J.R., & Jamieson, B.B. (2015). Human capital, social capital, and social network analysis: Implications for strategic human resource management. *Academy of Management Perspectives*, 29(3), 370-385.
- Karsh, B.-T., Waterson, P., & Holden, R.J. (2014). Crossing levels in systems ergonomics: A framework to support 'mesoergonomic inquiry'. *Applied Ergonomics*, 45(1), 45-54.
- Latour, B. (2005). *Reassembling the social: An introduction to actor-network-theory*. Oxford: Oxford University Press.
- Leonardi, P.M. (2012). Materiality, sociomateriality, and socio-technical systems: What do these terms mean? How are they different? Do we need them? In P.M. Leonardi, B.A. Nardi, & J. Kallinikos (Eds), *Materiality and organizing: Social interaction in a technological world* (pp. 25-48). Oxford: Oxford University Press.
- Neville, T., Salmon, P.M., Read, G.J.M., & Kalloniatis, A. (2016). Play on or call a foul? Testing and extending distributed situation awareness theory through sports officiating. *Theoretical Issues in Ergonomics Science*, 17(1), 80-103.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18, 87-127.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Orlikowski, W.J., & Scott, S.V. (2008). Sociomateriality: Challenging the separation of technology, work and organization. *The Academy of Management Annals*, 2(1), 433-474.
- Salmon, P.M, Walker, G., & Stanton, N.A. (2016). Pilot error versus sociotechnical systems failure: A distributed situation awareness analysis of Air France 447. *Theoretical Issues in Ergonomics Science*, 17(1), 64-79.
- Schraagen, J.M.C. (2015). Resilience and networks. *Proceedings of the Resilience Engineering Association 6th Symposium*, Lisbon, Portugal, June 22-25. Downloaded from: http://www.resilience-engineering-association.org/download/resources/symposium/symposium_2015/Schraagen_J.M.-Resilience-and-networks-Paper.pdf
- Siegel, A.W. (2017). Team reflection on weak resilience signals: Resilience enhancement of a rail sociotechnical system. PhD Thesis, University of Twente (CTIT Ph.D. Thesis Series No. 16-416). Enschede: Gildeprint.
- Siegel, A.W., & Schraagen, J.M.C. (2017a). Beyond procedures: Team reflection in a rail control centre to enhance resilience. *Safety Science*, 91, 181-191.
- Siegel, A.W., & Schraagen, J.M.C. (2017b). Team reflection makes resilience-related knowledge explicit through collaborative sensemaking: Observation study at a rail post. *Cognition, Technology & Work*, 19(1), 127-142.
- Simon, H.A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99-118.
- Simon, H.A. (1980). Cognitive science : The newest science of the artificial. *Cognitive Science*, 4, 33-46.
- Stanton, N.A., Stewart, R., Harris, D., Houghton, R.J., Baber, C., McMaster, R., Salmon, P.M., et al. (2006). Distributed situation awareness in dynamic systems: Theoretical development and application of an ergonomics methodology. *Ergonomics*, 49, 1288-1311.
- Woods, D.D. (2015). Four concepts for resilience and the implications for the future of resilience engineering. *Reliability Engineering and System Safety*, 141, 5-9.