21st International Command and Control Research and Technology Symposium 2016

C2 in a Complex Connected Battlespace

Topic 7: Methodological Development, Experimentation, Analysis, Assessment and Metrics

Testing AI watch-keepers in a mathematical model of networked OODA loops

Alexander C. Kalloniatis[1]

Defence Science and Technology Group, Department of Defence, 24 Scherger Drive, Canberra Airport, Australian Capital Territory, 2609, Australia

Over many years of conducting C2 studies at various command echelons I have often observed the use of operational HQ watch staff in the 'dead hours' to perform manual data processing work until operational imperatives occur. Is this the best use of human resources (given fatigue management best practice), and what is the impact in performance of the headquarters of replacing low rank military staff with smart, AI based systems that nevertheless still interact with senior officers? This paper presents a Use-Case of a mathematical model under development (presented at previous ICCRTS) based on networked oscillators as representative of operator OODA loops. The model is formulated as a set of coupled differential equations where nodes of the network may represent either human or non-human technological agents. I carry out an initial validation of the model using a threat scenario and dataset published at the 19th ICCRTS. Because empirical data on synchronisation was not sought in that study I use considerations from Contingency Theory as a proxy for measure of performance against a concrete threat scenario. I then model two interventions: firstly, by replacing human agents by smoother and faster (in terms of OODA loop) automated systems at the lower SA levels; secondly, by modelling an adaptive lagging mechanism that takes into account the heterogeneity of individual decision-making speeds. I discuss the improvements in performance of the headquarters in the various cases, but also the challenges of implementation.

Keywords: Headquarters, Network, OODA-loop

[1] Email: Alexander.Kalloniatis@dsto.defence.gov.au

1. Introduction

Eight years ago I proposed the development of a mathematical model for Command and Control (C2) that built on the key elements of cycles, timescales and structure as universal features regardless of the object of a C2 mission (Kalloniatis 2008). Since then, in two separate installments (Kalloniatis 2012, Kalloniatis and Zuparic 2014), I have shown how such a model behaves when describing adversarial C2 – the Blue-vs-Red model – and as the basis for modelling a Common Joint Staff System (CJSS) structured headquarters organisation. However each of these exercises was based on somewhat 'caricatured' organisational structures. In this paper I apply to the model real organisational data from a study reported in (Ali et al. 2014), validate the model, and use it explore the impact of Artificial Intelligence as decision aids in headquarters organisation.

The cyclic element of the model is based on John Boyd's Observe-Orient-Decide-Action (OODA) loop as a widely accepted concept for C2 (for recent scholarly analysis of the status of Boyd's model see (Hasik 2013) and (Osinga 2013)). However many other formalisms share this feature: the Endsley (1995, 2006) model for Situation Awareness (SA), in its broader sense, incorporates Decisions, Actions, and feedbacks[2]; Ullrich Neisser's (1976) Perception-Action process shares Boyd's fundamentally cyclic cognitive model; most military planning processes such as Australia's Joint Military Appreciation Process (JMAP), or the US based Operational Planning Process (OPP) reflect a similar pattern when placed in context of execution of plans and evaluation of the operation. Mathematically this is encoded as a limit cycle in a dynamical process. Time-scales reflect that the cycles will vary in frequency according to the position in the strategic-operational-tactical echelon, or may cut across all these levels. Finally, structure – organisational structure – provides the pattern of interactivity of agents in the C2 system, and is reflected in the model by networks drawing upon the mathematical tools of graph theory to describe them.

To these elements I add here one more associated with the fundamental component of C2, that which uniquely can exercise Command – the human (Pigeau and McCann 2002). Here, the human dimension is reflected in the individuality and precociousness of the individual that cannot be contained in any deterministic dynamical model. In mathematical modelling this is represented through stochasticity or 'noise'. Such effects, added to dynamical equations, seek to capture the *microscopically* unmanageable contribution of human factors. But rather than simply employ the common device of 'Gaussian' random fluctuations, I shall draw upon a new area of stochastic systems known as anomalous or 'fractional' diffusion (Metzler and Klafter 2000) built on Lévy processes where large random jumps may occur. This attempts to build into the model the propensity for experienced human decision-makers to anticipate solutions before completing a rational process (such as OODA, or JMAP or OPP). Such anticipation may be based on recognition of patterns in the problem from past experience (Klein 1998, Morewedge and Kahneman 2010).

With all these elements this paper genuinely seeks to offer insights into the benefits and limitations of technology in a connected C2 environment where the human person retains primacy. A particularly human face of the problem, manifest across the spectrum of command echelons, may be a restating of the third of the three 'Ds' - 'Dirty, Dangerous and Demeaning' - namely *Dull* work. I refer to the tasks often assigned to junior ranked watch-keeping staff in military headquarters, often in the dead of night

[2] Note that Grant and Kooter (2005) argue that 'Orient' in Boyd's OODA loop does not include planning whereas Endsley's model does; I argue that the projection into future states is simply not as elaborated with Boyd's level of fidelity and for the tempo of activity in his original context.

when day workers (who also often work punishing hours) are away, involving manual sifting through reams of data or documentation. The demoralising effect of such work on trained military personnel is that it is conducted under the harsh circumstances of shift-work (Dawson and Reid 1997) in an era where best-practice shift management recommends napping and maintenance of fitness where operations permit (Murphy 2002). In the age of smartphones and Google, can such work be assigned to Artificial Intelligence (AI) based autonomous systems liberating such personnel to tasks specifically requiring military skills and judgement? Note that to answer this question I do not attempt to model the AI in technological detail, but to capture their key attributes using the variables of the mathematical model. My analysis shows that, though upcoming technologies underpinned by AI may significantly improve the efficiency of a headquarters, the human context imposes limits to efficiency gains from technology.

Beyond this, my main goal in this paper is to offer a use-case as to how a Kuramoto-based model may be employed across the spectrum of C2 applications. Because the model is being developed under the scope of a research Fellowship, I do not have direct access to data tailored precisely for the model. Rather I re-use data from a previous study (Ali et al. 2014), that nevertheless richly documents the network structures for different operational contexts - and have made some calculated assumptions to fill in the gaps. The data itself draws upon two situations – a headquarters in steady-state mode and, alternatively, responding to a crisis requiring a rescue mission in a non-permissive environment. While the former is based on military personnel answering survey questions conducting real activity in a 72 hour period of life in the headquarters, the latter is based on the *same staff* responding to the *same survey questions* but for a hypothetical fusion of separate critical events ('search and rescue', and 'operations in non-permissive environments') with which the staff have deep experience. In this latter scenario, though there is here an 'adversary' to create the circumstances of non-permissiveness, the data concerns itself with the internals of the headquarters managing a response. In this respect, the model application develops the second of my previous studies of this type of mathematical model rather than the Blue-vs-Red context.

The Kuramoto model offers an advantage over solely computationally-defined agent-based models: its partial analytical tractability through fixed point and stability analysis (Kalloniatis 2010, Kalloniatis and Zuparic 2016) is valuable in model verification, deeper insight into model behaviours and connection with concepts from the mathematical literature on complex systems such as chaos, stability and entropy. Within the limitations of the data available, the present study provides a further stage in the validation of the model, after what may be termed the face-validity tests using simplified organisational structures in (Kalloniatis 2012, Kalloniatis and Zuparic 2014). Validation of models of social human systems (Sargent 1996) is a challenge at the best of times and is often pursued in the context of cross-validation across a spectrum of models such as from the typology of Rom Harré (1970): simple text descriptions and geometric diagrams (protomorphs), computational models of various sophistication (paramorphs) and human-in-the-loop representations (homeomorphs). In this paper I compare behaviours of a model based on differential equations (a paramorph) to text-based expectations from a field in organisational science known as Contingency Theory (a protomorph). The validation I conduct here may also be termed 'Event Validity' in the terminology of Sargent (1996).

The first half of the paper concerns itself with posing the model, explaining the dataset and associated operational contexts, tuning the model and finally validation. In the second half of the paper, I address the role of AI in a headquarters watch staff after explaining how such systems may be represented in a model of this form. The paper is completed with a discussion and conclusions.

2. The model – use case outline

Deterministic Model. The study of self-synchronisation in mathematically defined cooperative systems originates with Wiener (1961) and Winfree (1967) and has developed across the disciplines of mathematics, physics, biology and computation. Kuramoto (1984) distilled the elements into a first order differential equation whose network generalisation has become paradigmatic in the quantitative complex systems literature (Doerfler and Bullo 2014). The application of the Kuramoto model to robot-agent control, human-social and human-robot systems has precedent (Paley et al. 2007, Pluchino 2006, Mizumoto 2010). I use β_i to represent a time-dependent *phase* associated with node *i* of a complete network of *N* nodes, $\dot{\beta}_i$ is the angular rotation speed – or 'instantaneous frequency' - via the derivative of the phase with respect to time *t*, ω_i represents a 'natural' or 'intrinsic' frequency, usually randomly chosen from a statistical distribution, and σ is a coupling constant. The role of β_i as a phase is seen when it is inserted in the complex variable $\chi_i = e^{i\beta_i}$. A general network is represented using the adjacency matrix A_{ij} whose elements take value one if a link (or edge) exists between nodes *i* and *j* and zero otherwise. Though often presented for undirected graphs, the formalism can embrace the directed case where the index *j* labels all nodes that provide *inputs* to the node labelled by *i*, leading to a change of its state. The governing time evolution equation is then:

$$\dot{\beta}_i = \omega_i + \sigma \sum_{j=1}^{N} A_{ij} \sin(\beta_j - \beta_i) \,. \qquad (1)$$

The behaviour of the system can be visualised as points moving about the unit circle. For weak coupling the points will be randomly distributed around the circle, given the random individual frequencies. Above some threshold in coupling the points move with the same angular speed and tend to cluster ever closer together as the coupling is further increased. To measure the degree of synchronisation I adopt Kuramoto's (1984) order parameter:

$$r = \frac{1}{N} \left| \sum_{j=1}^{N} e^{i\beta_j} \right| . \qquad (2)$$

At strong coupling r smoothly converges to one ('complete synchronisation'), while at weak coupling it zig-zags around the value zero ('incoherence'). Synchronisation itself includes two cases: phase synchronisation is when all β_i are equal, giving r=1 exactly, while frequency synchronisation is when the difference between them is time-independent and may allow *r* to be far from value one but constant. The former can only occur exactly when frequencies are equal; the latter may approximate closer and closer to the former. By summing Eq.(1) over all nodes, it is straightforward to show that the average over the instantaneous frequencies $\dot{\beta}_i$ for the undirected network case equals the mean of the intrinsic frequencies; however, at exact phase synchrony even the directed network system rotates with this mean frequency. Special cases of the global order parameter *r* can be selected by only summing over specific sets of nodes, known as 'local order parameters', corresponding to specific staff teams, such as 'J2' separately from 'J3'. These will be indicated by appropriate subscripts.

Mapping to C2. Translating this to C2, to reiterate the mapping established in (Kalloniatis 2008), the phase $\beta_i(t)$ represents the point in a continuous decision (or OODA) cycle of agent i at some time *t*.

The network represents the C2 structure itself: the relationships of agents that provide inputs for mutual adjustment of individual decision cycles.

The coupling σ is a measure of interactivity between C2 nodes. The scholarship around coupling in organisational theory and C2 includes Weick (1976), Perrow (1984), Beekun and Ginn (2001a,b), Lloyd, Markham and Dodd (2006), and Stanton (2006). Weick, referring to 'loose coupling' (in contrast to 'tight') as advantageous to an organisation, lists some 15 notions of coupling which may be aggregated into two exclusive notions: coupling as connectivity, and coupling as intensity. Most attempts to measure coupling have devolved to the former using network metrics (Beekun and Ginn 2001a,b) rather than strength, intensity or frequency of interaction on the network links. In my approach, coupling may empirically refer to the speed of change in decision state by one node in response to a change of decision state by a connected partner or adversary.

The 2π-periodicity of the sine function is appropriate in that it locally synchronises decision cycles within the 'current phase'. The frequency ω_i is how many decision cycles per unit time can be achieved by agent i. This is chosen from a random distribution, representing the underlying heterogeneity between individual decision makers in the C2 system. Training and discipline can narrow that distribution; namely, introducing more homogeneity in the population of decision makers. But the intent is nonetheless to retain some degree of heterogeneity. Moreover, one does not have the luxury of 'managing' that heterogeneity: the C2 system is not designed with individuals of certain frequencies placed deliberately at certain nodes.

The Kuramoto model is ultimately a model quantifying *self-synchronisation* on networks. Certainly in the NCW literature, such as Alberts and Hayes (2007) and references therein, the desired self-synchronisation is applied to *activity* in the external environment. My proposition is that the precursor to this is synchronisation of *decision cycles* and therein mapping the phase of the Kuramoto model to the decision cycle; another implementation of the Kuramoto model is possible at the level of activity and is that used in (Dekker 2007, 2011). These two options are not very far apart: a decision cycle in a context such as a headquarters will very often leave a trail of external artefacts - published or draft documents, emails, chat or verbal communication - that indicate the stage of OODA of a unit or individual; these artefacts are thus points of reference for another in the same organisation in synchronising their cycle. In other words, even the cognitive stages of Observe-Orient-Decide involve some form of social enterprise, when one steps beyond Boyd's original application to the isolated fighter pilot alone in the cockpit.

The reference to artefacts hints at a further refinement of the adoption of Kuramoto to C2 which I adopt in this paper, that information objects constitute nodes in the network as well. This is not novel in C2 models, figuring heavily in the Distributed Situation Awareness (DSA) approach of Stanton (2006) and a related method to which I contributed, the application of which several years ago lead to the data I employ here (Ali et al. 2014), namely the Situation Awareness Weighted Network (SAWN).

Applying Noise. The deterministic Kuramoto system is transformed into a particular stochastic differential equation (or 'Langevin equation') by adding terms $\Gamma_i(t)$ to Eq.(1) that are erratically fluctuating functions of time (Khoshbakht, Shahbazi, Samani 2008). Often Gaussian White Noise is employed, where the function values are uncorrelated from instant to instant, have zero mean and some fixed variance. In this paper, at each time-step I draw the values of the function from a distribution that follows a power-law with index α, $0<\alpha\leq 2$, known as the stable Lévy law (Khinchine and Lévy 1936, Lévy 1937): the probability of a step in space of phase of size x is $1/|x|^{(1+\alpha)}$ (for large

x) for a random walk in one dimension (Metzler and Klafter 2000). Such distributions admit 'fat' or 'heavy tails' and approach the Gaussian case when $\alpha \to 2$. The fat tails imply that occasionally large jumps may occur. The parameter α controls both the size and frequency of large jumps in the random walk model. Values close to two give very few and very small 'large' jumps; values less than one give very frequent and large jumps to such an extent that for $\alpha \leq 1$ the Lévy process has no finite mean[3]. The additive noise will be applied to individual nodes as a time-dependent addition to the frequencies ω_i. A more rigorous study of the Kuramoto model on theoretical networks with such noise was recently undertaken by Kalloniatis and Roberts (2016).

Noise in C2 systems. Such Lévy noise for the C2 context is appropriate on two grounds. Firstly, noise models the degree to which the human dimension cannot be microscopically tracked and modelled, analogous to the microscopic collisions of suspended colloid particles in the original observation of Brownian motion. Additive noise specifically means there is a random time-dependent element to how fast a cycle is completed: no individual processes information or makes decisions with the same speed in every instance as individual internal factors (such as, for example, mood, health, confusion) may vary from instance to instance. That elementary binary human (and animal) decisions may be subject to Gaussian noise is well known from the Ratcliff model in cognitive psychology – for a review see (Ratcliff and McKoon 2008). However, in complex decision making under uncertainty 'experts' are known to jump (and backtrack) according to a variety of less quantitative models, such as Klein's Recognition Priming (Klein 1998) or in the context of Wicked Problems (Rittel and Webber 1973). *Lévy noise seeks to quantify such models*, though there is no data available of the quality used in the Ratcliff model.

Frustrations. One more addition may be made to the model that is known in the literature: shifts φ may be added inside the sine function, $\beta_j - \beta_i \to \beta_j - \beta_i + \varphi_j$, leading to a variation known as the frustrated Kuramoto model (Kirkland and Severini 2015) or Sakaguchi-Kuramoto (1986) model. This has the effect of relaxing the interaction from seeking to phase synchronise to merely frequency synchronising. For C2 applications such frustrations may build in the property that two agents in the C2 network do not need to precisely synchronise decision loops, but they may be shifted with respect to each other. This enables, in addition to the network structure, well-known C2 processes to be incorporated in the model; for example, logistics planners may need to be staggered in the planning cycle in relation to operational planners or intelligence estimators. I forego this modelling aspect here simply to allow the interactions governed by the network structure to dynamically find the degree of synchronisation *possible* in the system, rather than forcing a sequence by explicit frustration parameters[4]. However I return to this idea toward the end of the paper.

3. The Scenarios

To populate this model with data I draw upon an empirical study conducted with colleagues, and reported in 2014, within an Australian Defence Force (ADF) military headquarters which interacts with a number of higher and subordinate Joint Task Force (JTF) or Task Group headquarters and single service units. The headquarters may be deployed in a particular theatre or based in Australia; the ADF conducts a range of operations in its region and globally both individually and in coalition.

[3] Many will be aware of the Central Limit Theorem that a large number of iterates of a random variable, *each with well-defined mean and variance*, will approach the normal distribution. The stable Lévy law generalises this for when the variance is no longer kept finite (Lévy 1936).
[4] The 'C2-process' may thus be deemed to be *emergent*, a consequence of intrinsic dynamical interactions.

More significant is that these operations require both air and maritime platforms requiring a joint response. What potentially renders conditions in these areas as non-permissive are threats that may range from actions of nation states to non-state based armed groups such as terrorist networks or pirates. The key staff in the headquarters provides ongoing information and analysis of events and initial coordination of responses for specific areas of operation. To this end, they rely on information systems and displays, products, services from other organisations, and personal interactions. Staff include regular 'day-workers' and shift watch staff on a 24 hour/7 day roster.

The headquarters is structured using the Common/Continental Joint Staff System (CJSS). Within this naming system for headquarters functions are the 'J3' operations staff, who control the movements of tactical units in areas of operations, and the 'J2' intelligence staff who monitor possible threats in the same or adjacent areas. These same staff issue routine briefs, drawing from a Common Operating Picture (COP) that they themselves maintain. The J3 watch staff are also responsible for initial responses to events outside formal standing JTF or Task Group Areas of Operation. Elsewhere in the J2, analysts provide support with deeper examination of events.

As described in Ali et al. (2014), and summarised below, data was collected on how staff interacted with each other directly or via information objects and displays to build their own or another's SA in the context of two scenarios. *Routine* means the steady battle-rhythm of the headquarters built on the key component of the daily Commander's brief in the morning informed by ongoing Intelligence, Surveillance and Reconnaissance (ISR) operations, and shift handovers every 12 hours. Observations and surveys were conducted *in situ* in a period in 2014 where no exercises or other specific activities were planned. Fortunately, no unanticipated events occurred during this time permitting an appropriate base-line to be established. *Crisis* referred to an event involving a hypothetical search-and-rescue operation by an Australian platform that happens to be close to the area of an incident but in a context where a potential threat was emergent. To this end, staff in J2 and J3 in the headquarters respectively would conduct Intelligence Preparation of the Battlespace (IPB) and Crisis Action Planning (CAP). Such staff processes are similar across coalition nations, with details readily accessible through openly published US Department of Defence Joint Publications and Field Manuals.

4. The Data

SAWN approach. The data describing headquarters response activities to these two scenarios was collected in the context of development and application of the SAWN methodology which seeks to unify the two prevailing models of SA, Endsley's (1995) individual SA model of three levels – Perception, Comprehension and Projection - and the aforementioned DSA approach. SAWN, like DSA, employs the network paradigm – both social and information networks[5]. The Endsley (1995, 2006) model supplements this by weighting links in the network according to the level of SA intended or sought by agents at the nodes of the network. DSA purists would object to these levels, but these attributes may be viewed more appropriately as discrete increments in the cyclic decision process, such as OODA. Indeed the study (Ali et al. 2014) demonstrated both the non-monotonicity in the progress of individuals through these levels and the multiplicity of views of the SA of interacting agents.

The instrument for the data collection was a survey of two sets of 12 questions representing a finer granulation of the three levels and covering such activities as having attention drawn to an asset in the operational environment, identifying its character, its relationship to other assets, its current and future

[5] DSA also uses task models but in SAWN only the social and information networks are combined in applications to date.

actions within a prescribed time-horizon. The two sets correspond to SA associated with Products (Information Objects and displays) and Organisations (people), and further break down into questions how one interacts for someone else's information – 'Pushing' – and for one's own – 'Pulling'. The staff that the study was directly able to sample ranked from the O5 level down to enlisted personnel. After some aggregation, the set of nodes is reproduced here (from Ali et al. 2014) in Table 1.

Table 1 Nodes in the network after aggregation of raw data; italics indicate organisational nodes that were directly sampled. Where a 'W' appears in the name, the section or individuals are Shift Workers, with the exception of OJTF that has its own watch and therefore is also treated as shift based. Taken from (Ali et al. 2014).

Organisations	Characterisation	Products	Characterisation
OCmnd	Commander, or function/section Heads	Pemail	Formal email that documents information (rather than ad hoc communication)
		Pweb	Documents acquired through intranet
		Pbrief	A fused brief from J3 and J2 staff (see below)
		PCOP	Parts of a COP maintained by J3 and J2 staff
OJ3W1	Senior Officers in J3 Watch	PJ3Brief	Routine Briefs from J3 Watch staff
OJ3W2	Junior Officers/NCOs in J3 Watch	PJ3Ord	Orders issued by J3 Watch staff
OJ3WS	Support officers in J3 Watch		
OJ3S	Support officers in J3	PRpt	Formal Reports by military staff
OJ2W1	Senior Officers in J2 Watch	PJ2WBrief	Routine Briefs by J2 Watch staff
OJ2W2	Junior Officers/NCOs in J2 Watch	Padhoc	Ad hoc document or PowerPoint slide to capture ongoing events
OJ2A	Analysts in J2 Staff	PJ2AnBrief	Analysis Brief by J2 staff
		PThrAss	Threat Assessments
OJ2S	Support officers/NCOs in J2 staff	Popen	Open source information – internet
OStrat	Strategic Level Staff external to the HQ	Penv	Reports on physical environmental conditions
OExt	Organisations external to Australia	PExt	Reports from organisations external to Australia
OInt	Other intelligence organisations	PInt	Reports from other intelligence organisations
OJTF	JTF or TG HQs	PJTF	Brief or formal signal from JTF or TG HQ
OSS1	Single Service (Army, Navy, Air Force)	Pdef	Signals from Single Service units
OWOG	Non-defence government departments	PWOG	Reports from non-defence government departments

For steady-state conditions, staff completed the survey while on-duty reflecting on their current activities. For the crisis scenario, a separate task model was developed in a workshop of SMEs of representatives from across the headquarters, including planning (J5) and single service staff. A condensation of this model was presented to the original survey participants who were then asked to answer the questions in view of the tasks detailed for the crisis response. Further details are given in Ali et al. (2014).

Networks. The net result of this data collection was a set of four networks (Ali et al. 2014), two each representing Push and Pull for steady-state conditions, and two for the Crisis. Each network is a semi-bipartite weighted digraph: *semi*-bipartite because information objects do not of their own accord interact with each other and thus no links exists between such nodes whereas humans do and so nodes in the pure 'social network' are linked.

For this paper the weights are of lesser importance – and so are set to one if non-zero. Also not modelled directly here is whether the interaction is pushed or pulled. The directionality is, however, important in that the recipient of information (whether sought by them or sent to them 'unsolicited' by another) will undergo a change of decision state. By standardizing the adjacency matrices in this way, for example having rows labelled by *i* representing the source and columns *j* representing the sink, push and pull matrices may simply be added:

$$A_{ij} = \frac{1}{2}\left\lceil A_{ij}^{(pull)} + A_{ij}^{(push)} \right\rceil \quad ,$$

where the ceiling function $\lceil\rceil$ means that the final entries in A will be one or zero. The resulting two networks, for steady-state and crisis scenarios, are shown in the Appendix, in Figure 11. Note that they are shown there without arrow-heads (namely, 'undirected') simply for clarity – in computations I shall always be dealing with directed networks in the following.

Natural Frequencies and Coupling. The concern of the study in (Ali et al., 2014) covered neither such things as decision speeds nor coupling strength, and therefore I have no available data on speeds or frequencies of OODA loops. Indeed, despite its provenance in military C2, there is very little direct measurement of such things. What is known for this case is that shift-workers work for 12 hours and day workers approximately 10-12 hour days. I do not attempt to model individual shift-workers, rather their positions. Thus shift-based positions are available 24/7. Assuming that, within a uniform error range, an operator has the same OODA frequency as number of decisions/unit time regardless of whether their work is shift or day based, I may simply compute the ratio of day-based position hours to shift based giving 0.5. Thus I assign ω=1 for shift-workers and ω=0.5 for day-workers. This leaves an overall absolute scale factor unaccounted for. However, the same scale applies to the coupling strength. As will be explained, the coupling will be obtained by calibrating the model to a certain level of performance at steady-state. This serves then to fix the unknown scale in the frequencies.

Of course, information artefacts – at least those in the headquarters of our study – do not engage *of themselves* in a decision cycle. Therefore the natural frequency of these is exactly zero. However, in the model they do update their OODA state by human operators pushing information to them. Thus, a dynamical oscillator phase may be assigned to them. This element of the model is quite apposite for modelling the *burden* of headquarters staff in maintaining a range of information products: they do not update themselves, therefore the requirement that the entire system of people and information achieve a level of synchronisation is effortful – and the model reflects this. By way of contrast, I observe here the opportunity for modelling AI by considering 'smart' information objects that indeed do engage in a self-driven OODA loop, as has been proposed for modelling autonomous systems by Proud et al. (2003). Thus, AI artefacts offer the scope for relieving the burden of information work – a hypothesis that will be tested in this paper.

Returning to the 'dumb' information objects in our study, some of these – such as material freely available on the internet, 'P_{open}' – are created and updated by individuals *beyond our sample*, and therefore can never in the present model 'update' their state. In solving the Kuramoto system such phase variables remain static and coupling the remainder of the nodes to them has the effect of freezing the entire system, a prodigious modelling artefact. I therefore decouple such nodes when solving the Kuramoto dynamics.

To account for variability of individual human decision speeds, I draw frequencies from uniform distributions within the deviations [-0.1,0.1] of the above frequencies. The narrowness of this distribution reflects that the staff in the study are experienced military who have worked together for at least 6 months in that posting (and many in previous postings). The results I report do not depend sensitively on the specific numerical value of this range. What is important is that the distributions do not overlap unless day workers begin to extend their working hours, dangerous for quality of decision-making as studies of fatigue are able to relate performance levels for numbers of hours-awake to an equivalent blood-alcohol level (Dawson and Reid 1997). Figure 1 shows the range of frequency choices across shift and day workers and 'dumb' information objects for the numerical computations presented later.

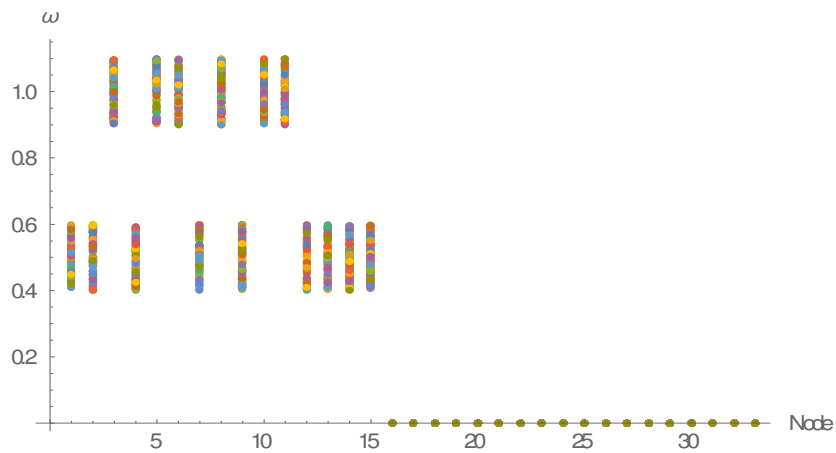

Figure 1 Sampling of the frequencies for 100 instances, each coloured dot representing a specific instance, with node index on the horizontal axis and the frequency choice on the vertical; shift-workers have their frequency centred about ω=1 drawing on a uniform distribution, day workers about ω=0.5, and information artefacts have ω=0.

Noise. Again, there is an absence of data here but significant insight available from qualitative studies of military decision-making to inform assumptions. For steady-state, because of the routine nature of tasks, I employ Gaussian noise. The Gaussian based Ratcliff model, though critical for motivating noise in my model, is concerned with quite simplistic decisions and cannot be readily adapted to the context of cyclic decisions (Ratcliff and McKoon 2008). Instead I choose zero mean noise – the native frequencies provide for the drift here – and a variance of the noise that provides for its support (the region over which the probability density function is appreciably greater than zero) well within a quarter-of-a-cycle (π/2); in other words, noisy decision-making for experienced operators in routine conditions has fluctuations within an OODA 'step'. Numerical trials lead to a noise variance $D = 2\sqrt{2}(0.2)^2$ employed in a density for the probability of drawing value x of the form $\exp(-x^2/D^2)/\sqrt{\pi D}$.

For the crisis scenario, given the increased uncertainty and the the likelihood that experienced staff may apply informed intuition (Klein 1998) I use, as mentioned, the stable Lévy distribution because of its heavy tails. One parameter of such distributions matches the variance D in the limit α=2, so I maintain the value of D from the Gaussian case. Stable Lévy noise may also be *skewed,* namely large jumps may only occur in one direction while Gaussian small incremental fluctuations may apply in the other. This is reasonably applied to seasoned decision-makers, particularly in a crisis in order to maintain, forward momentum in decisions. I therefore apply a maximal bias in the underlying random noise model. With these two settings of the noise I thereby represent the environment in terms of the uncertainty it generates in individual decision-making.

The choice of α is fundamental. I use α=1.4 as typical of a value with infrequent jumps of the order of half a cycle, π, or more. Observe in the blue curve of the left hand plot of Figure 2 that α=1.4 is approximately the value at which the probability that a jump exceeds π, based on a sample of 100, becomes non-zero. By inspection of the plot, the results do not depend sensitively on this choice. In the numerical solution of the Kuramoto model this is implemented, for example by cutting off very large jumps at, say, π (note also the jump at α=1 in the figure, indicating the pathological case of this choice – known as the Cauchy distribution). Based on this, the remaining plots of Figure 2 show an example of the probability density function (middle) and time-series (right) for this choice of α, contrasting with Gaussian noise of the same strength (used for the steady-state scenario, with parameter *D* as discussed earlier). Observe the fat-tail in the positive direction for the blue curve in the middle plot and the occasional large positive jumps in the time-series, characteristic of this distribution.

5. Measuring and judging performance

In the original SAWN study (Ali et al. 2014), the success in performance of the headquarters staff in their work – either maintaining steady-state or responding to the crisis – was not measured. The former may be taken as a given – indeed the broader aspects of a preceding study conducted by the authors of (Ali et al. 2014) involved interviews with senior officers in the headquarters ascertaining satisfaction with the performance of the watch-keeping staff under these circumstances.

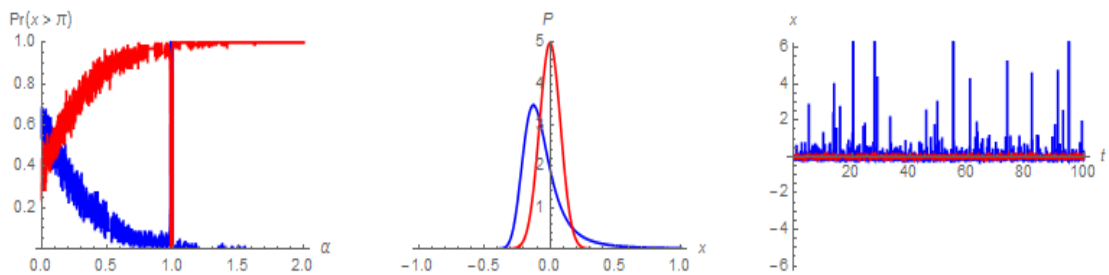

Figure 2 Representations of the stable Lévy distribution. Left: the probability that a random choice from 100 samples of a maximally skewed stable Lévy distribution exceeds π (blue) or is less than π (red) as a function of the index α. Middle: The probability density function for stable Lévy distribution for α=1.4 (blue) and the Gaussian distribution (red). Right: a time series drawing from the stable Lévy distribution for α=1.4 (blue) and the Gaussian distribution (red).

The Kuramoto model provides a natural measure of performance, the order parameter *r* of Eq.(2), though the model itself does not assume that synchronisation is good or bad. In applying this to headquarters staff, *frequency synchronisation will be taken as a proxy measure of performance*. Staff need not be at the identically same point in their respective OODA cycles – indeed given the time to process information at one level of SA and add-value to it to bring it to another (as the original SAWN study verified in Ali et al. (2014)), it is impossible to expect two staff members to be at the same OODA point. However, the difference in different members' OODA states should be constant in time to reflect a team performing coherently. I therefore apply Measures of Performance like Eq.(2) to the three groupings:

- The J3 staff (watch and day workers): OJ3W1, OJ3W2, OJ3WS, OJ3S;
- The J2 staff (watch and day workers): OJ2W1, OJ2W3, OJ2A, OJ2S;
- The combined headquarters Watch: OJ3W1, OJ3W2, OJ3WS, OJ2W1, OJ2W2.

The order parameters for each of these are then, with appropriate normalisations,

$$r_{J3} = \frac{1}{4}\left|\sum_{j\in J3} e^{i\beta_j}\right|, \quad r_{J2} = \frac{1}{4}\left|\sum_{j\in J2} e^{i\beta_j}\right|, \quad r_{Watch} = \frac{1}{5}\left|\sum_{j\in Watch} e^{i\beta_j}\right|. \quad (3)$$

When these become time-independent – the system dynamically reaches 'equilibrium' - the appropriate level of performance will be deemed to have been met.

What remains unmeasured in the data is the performance in response to the crisis scenario. A key observation made in Ali et al. (2014) – and visible in the networks shown in Figure 11 – is the increased number of links in the crisis scenario compared to steady-state. This behavior is well-known in the literature as 'uncertainty reduction theory', which posits that people engage in intensified information seeking activity to alleviate uncertainty when they are under threat or facing a crisis, perceived or actual (Berger and Calabrese 1975, Afifi and Weiner 2002). Overall, scholars have shown that in conditions of high uncertainty and task complexity managers seek more information from different sources for decision making than for routine tasks and that they show preference for verbal as opposed to written media (de Alwis, Majid, and Chaudry 2006). What remained unsaid was that this information-seeking potentially exposes the senior watch staff to information deluge and decision paralysis. Quantitatively, the key differences between the networks for the two scenarios lie in the increased degree, by the amount 4 and 2 respectively, of the OJ3W1 and OJ2W1 nodes, the highest (military) ranked staff of the sample. Therefore the network in the crisis scenario is *more* centralised.

Contingency Theory (CT) provides a prediction for this circumstance. Largely a qualitative approach (Mintzberg 1979, Donaldson 2001), CT posits that organisational structures must be fit-for-purpose for the contingencies they confront: routine, predictable environments are suited for hierarchical/mechanistic organisational structures while dynamic, uncertain environments require organic networks. On this basis, the data for the crisis scenario suggests the headquarters is exposed to some risk – of information overload of the central decision-makers (not measurable in the model) and of imbalance in the relationships between staff being drawn between many information sources and centralised decision-makers (which is measurable). In applying the crisis data to the model, I will test whether there is an *echo* of such behaviour constituting the 'event' for the validation.

6. Tuning the model

I calibrate the model to the steady-state scenario. Having now specified the data for the network and having made judicious assumptions about the frequencies and noise the only remaining parameter to be tuned is the coupling constant, σ. This tuning is as follows. I start with coupling set to zero and numerically solve the Kuramoto system for the steady-state data for an ensemble of 100 instances of random initial conditions, frequency and noise choices using `Mathematica.` I increment the coupling to larger values and re-solve the system. For each solution set, I compute over the ensemble for each of the local order parameters the mean, first and third quartiles and maximum/minimum values at each time step. At each coupling value the same random seed is applied. For each coupling choice the computation requires of the order of 3 minutes on a standard laptop or desktop machine. I increment the coupling until the mean of each order parameters show behaviour over time converging to a constant.

In Figure 3 a box-whisker[6] plot of the computational output for the steady-state scenario is shown for the three local order parameters as functions of time up to time t=100, ignoring a transient before t=10. Three values of coupling are shown in the three rows: at zero coupling and the coupling just below the point of synchronisation, σ=0.4, and coupling at synchrony for all parameters, σ=0.6. The particular behavior in the middle row of plots reflects that, whereas the majority of the ensemble show reasonably good synchronisation, a number involve fragmentation of a small subset of agents who are periodically are able to track the rest of their unit, but then break away and must skip a loop or jump ahead an entire loop to rejoin the 'pack'. Each instance of these may have a slightly different period, giving the nearly periodic pattern for the maxima and minima.

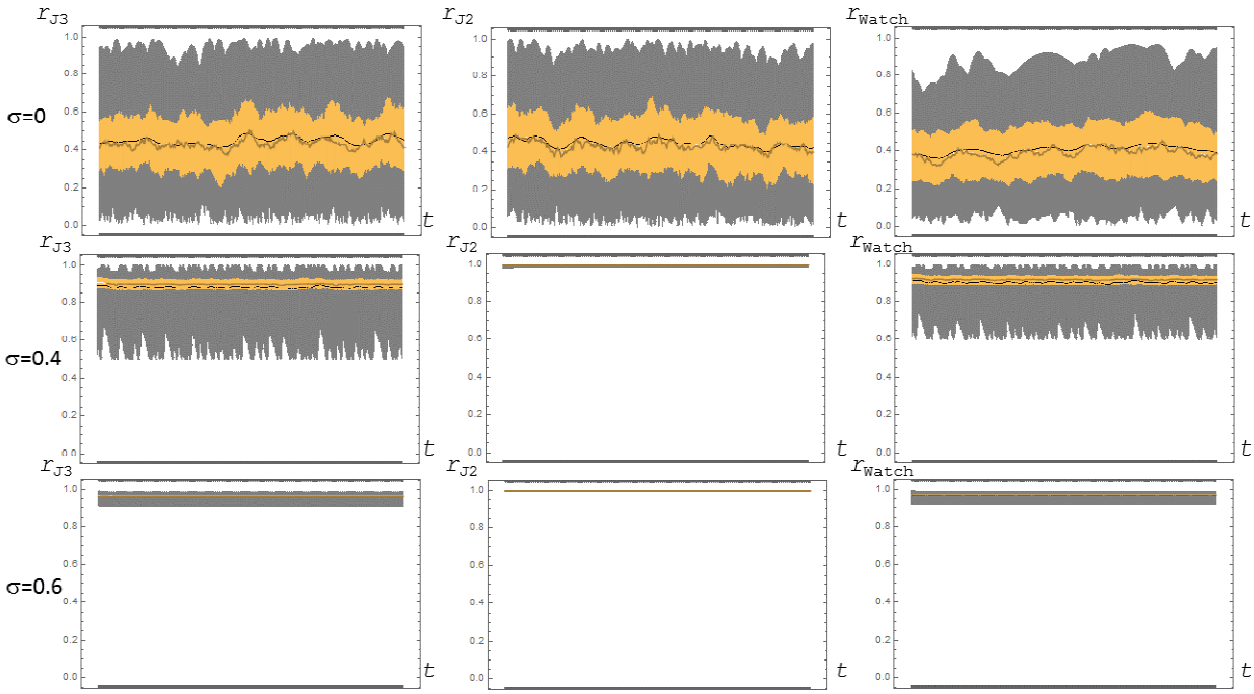

Figure 3 Box-whisker plots of local order parameters Eqs.(3) for steady-state SAWN scenario data at coupling σ=0 (top row), σ=0.4 (middle row) and σ=0.6 (bottom row) as functions of time. Each point represents statistics over an ensemble of 100 runs at a point in time of the dynamics by solving the Kuramoto model with SAWN steady-state scenario data applied. Grey indicates the extension of the data between minimum and maximum values, orange enclose the quartiles and light grey line indicates the median (mid-point of the quartile bars), while the solid line shows the mean.

At σ=0.6 (last row of Figure 3) all order parameters 'flat line', but with different variance for the J3 and J2. Overall, I observe that J2 synchronises more readily and closer to phase synchrony than J3; correspondingly the whole watch is somewhat less well synchronised than the J2 sections of it. Thereby, I fix the coupling to σ=0.6. Arguably, the coupling could be higher and still satisfy the requirement of coherent J3, J2 and Watch teams. However, in the real world, coupling is an effortful process and so a collection of individuals may evolve their interactions to the lowest coupling manageable to achieve outcomes, reflecting a Principle of Least Effort (as in Zipf's Law for word usage in language). Observe that once this regime is achieved, the degree of synchronisation is in fact quite high, above 0.9 generally. However, this is not exact *phase* synchrony in any of the groups –

[6] A box-whisker plot combines the first and third quartiles of a data set and the median, giving the 'box', as well as the minimum and maximum values, giving the 'whiskers'; the median, as the 50% quartile, wll always be at the mid-point of the box. The mean of the data may also be plotted.

inspection of the raw numerical data shows a spread of individual phases over all instances and times within a range (−π/4,π/4). Note that in the absence of Gaussian noise the critical point is at σ=0.4.

One might think that at the critical coupling all is as might be hoped for the staff. As an illustration of the impact of the production of information artefacts in the dynamics of the system, I plot for a specific instance of solution of the model at critical coupling in Figure 4 the instantaneous frequency of the staff (blue) and the information products (red) comparing them to the natural frequency of the staff (grey). Recall that for the directed network case, as I have here, the system in general will not synchronise to the mean of the intrinsic frequencies; however the closer to phase synchrony the more the mean intrinsic frequency dominates the aggregate behavior of the system. Therefore, as a consequence of the interactions, particularly with information artefacts (the producing and use of them) that the actual instantaneous frequency of staff is less than their associated intrinsic frequency – both for day and shift-workers. The *effective* rate of decision-making is slowed by the effort – one might call it, the *inertia* – of producing 'dumb' information artefacts.

This behavior cannot be improved by increasing the coupling which only serves to bring to convergence the total population of staff and products to the mean of the entire ensemble of frequencies; this will always be less than that of the slowest human decision-makers while the products have no ability to self-update their OODA state. I return to this later in the paper.

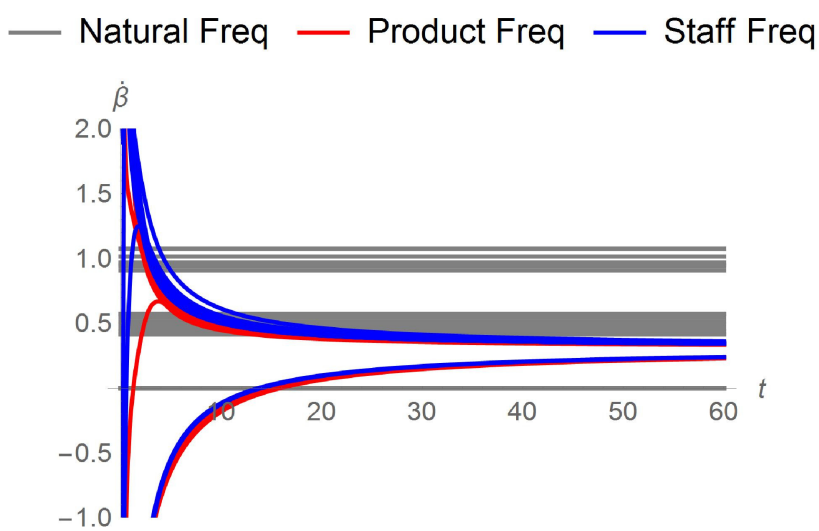

Figure 4 Instantaneous frequencies for a specific instance of the 100 solutions of the model for steady-state scenario data at the critical coupling showing the natural frequency (grey), the product frequency (red) and the staff frequency (blue)

7. Validation of the model

I now apply to the Kuramoto model the SAWN crisis scenario network data and the skewed Lévy distribution noise using the coupling σ=0.6 as calibrated to the steady-state conditions. In applying the same coupling I make the assumption that given the same team and same means of communication an unexpected crisis will not provide immediate opportunity for coupling to be tightened – again the Principle of Least Effort. Put alternately, operators who have adapted to a certain rate of responding to each other will not readily change it except when subject to some feedback – namely that they are failing to adequately respond to the crisis coherently, evidenced, for example in people backtracking or rushing ahead to recover from mistimed processes. I solve the Kuramoto model again for 100

instances, again using the same random seed as for the steady-state scenario. The resulting local order parameters are plotted in the box-whisker plots of Figure 5.

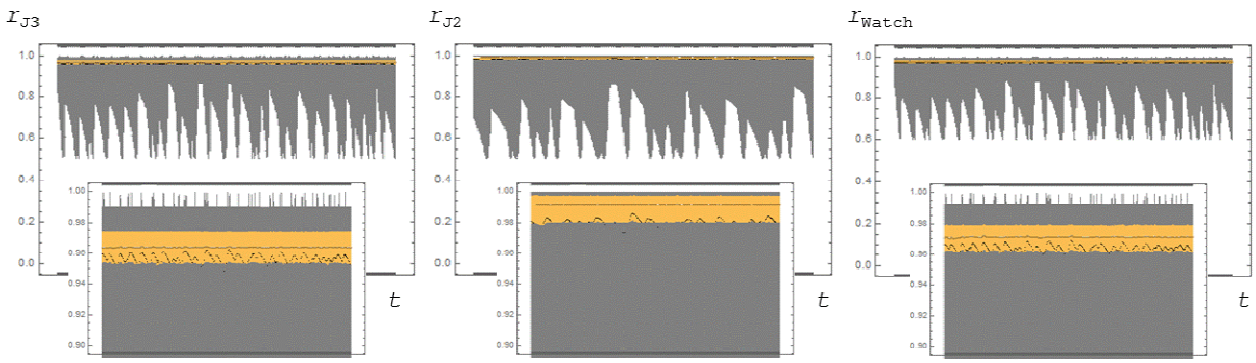

Figure 5 Box-whisker plots of local order parameters Eqs.(3) for crisis SAWN scenario data at coupling σ=0.6 (top row) and zoom in to the range (0.8,1) inset below. Colour scheme as in Figure 3.

Observe now that a degree of incoherence occurs under these conditions. Specifically, the quartiles are reasonably constant, and therefore so too is the median. However the mean, seen in the fluctuating dotted line in the zoom-in of the plots (insets zoomed into the range $0.9 \leq r \leq 1$ below the main three plots in Figure 5), is clearly time-dependent. The maximum and minimum values also show significant fluctuations, showing that *some instances fail to synchronise.* Thus there is a non-vanishing probability of failure. The incoherence is most severe in the J3 unit, but, by virtue of the same degree of fluctuation evident in the Watch behavior, this is not purely because of a disjointedness between day-workers and shift-workers but lies in jumping and back-tracking of staff within the watch itself. The mean over all ensemble data at each time step shows a distinct periodicity consistent with a distinct set of agents in the network managing to synchronise for some period of time but then dislocating from the main group and slowing down or speeding up to rejoin the collective decision cycle.

Synchronisation in this scenario is restored by doubling the coupling, from σ=0.6 το σ=1.2. In the real world this corresponds to individuals speeding up – doubling - their responsiveness to each other. Naturally, military professionals are capable of this. This model cannot answer – with present data or even with assumptions – how long this might take. The probabilistic nature of this model means that there is a risk that an error may be made by the time this adjustment occurs. Inspection of the individual solutions reveals 13 cases of the 100 instances failing to synchronise, namely probability of failure of 13%. Some of these only involve external entities such as JTFs/TGs or the Strategic level (no less concerning). Within the watch, the nodes that typically fail are the subordinate ranked ones, responsible nevertheless, from the original SAWN data, for Perception data for the building of SA. Naturally, regular training and exposure to crises enables a team to make adjustments in coupling (or network) more quickly. Indeed, there may even be a hysteresis effect in that after crises teams may relax to a steady-state value of coupling that is higher than before the crisis. All of these aspects are beyond the capability of a model such as this.

The source of this higher required coupling lies in the increased hierarchy of the network for the crisis scenario; tree graphs are known to synchronise more poorly than more connected graphs (Kalloniatis 2014). This is consistent with the source of risk for hierarchies in dynamic uncertain contexts from Contingency Theory. This match between model behavior and empirical expectations constitutes an

‘event-based’ (Sargent 1996) validation of the Kuramoto model applied to C2. Note that if the steady-state scenario network were subject to the same form of stable Lévy noise the critical coupling would have been σ=1, still below that required for the crisis scenario network to achieve synchrony. Therefore, the incoherence is not purely a consequence of the noise choice but also the relative lack of robustness of the network structure.

Having validated, to this degree, the model, its power now lies in its ability to predict changes in behaviour under *a large range of interventions* in the C2 arrangements for the crisis scenario. Three are to locally redesign the communication networks such that appropriate sources of information are accessed without further centralizing the network, or to apply collective or individual training to narrow the distribution range of decision-making speeds, or tighten the responsiveness between any two individuals. In this paper I take another path, given the focus on AI and autonomous systems.

8. Testing AI-based interventions

I have thus far alluded numerous times to the property of even the most modern of information objects in a military headquarters (in my experience) to require human input for state-updates – they do not have a self-driving OODA loop, and often the updating of this state is the mundane responsibility of the most junior ranked military staff on watch. Both the information artefacts and the junior staff offer an opportunity for the replacement by AI-based decision aids. For the former, the key to incorporation is the assignment of a native frequency to the information objects represented in the SAWN data. I also assume that such entities will be ‘noise-free’ in the spirit of how I use noise to model human variability and intuition. For the latter, I represent the replacement of the junior watch staff with (again) noise-free agents. I continue to use the critical coupling determined from the steady-state scenario but apply these changes to the crisis scenario to see how such interventions may improve the observed loss of coherence here.

Smart Information Objects. I set here the frequency of the information objects exactly to one, the same as the fastest (watch-keeping) staff but with no noise. The usual box-whisker plot for order parameters is shown in Figure 6 which should be compared to Figure 5. Synchronisation has improved, but is not at the level of the steady-state case at the same coupling; there remain instances where the system fails to frequency synchronise. On the positive side, the instantaneous frequencies of the staff now approach ω=1, in contrast to Figure 4.

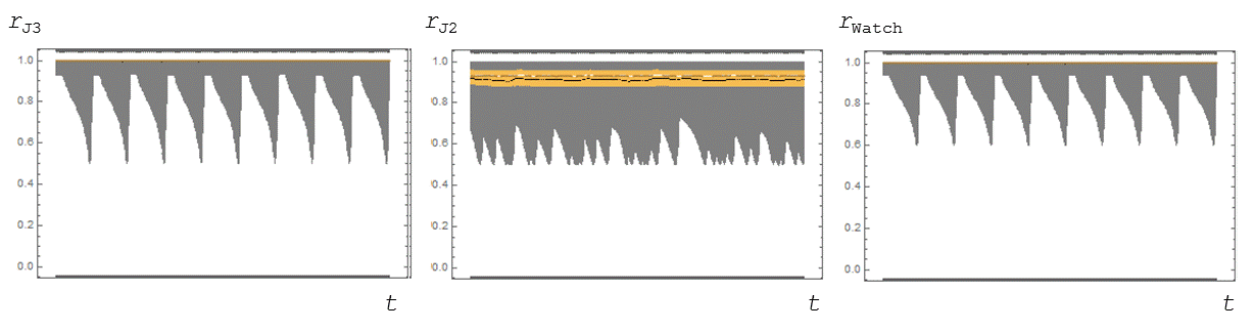

Figure 6 Box-whisker plots of the local order parameters Eqs.(3) with the SAWN crisis data at σ=0.6 but natural frequencies of all information objects set to 1. Colour scheme as in Figure 3.

Naturally increasing the coupling will improve the synchronisation. It may seem plausible that increasing the frequency of the information object OODA loops will help – faster AI! In fact this

makes things worse, as can be seen in the order parameter for the J2 staff who interact with more information systems and therefore fall into permanent periodic behaviour, as evident in Figure 7.

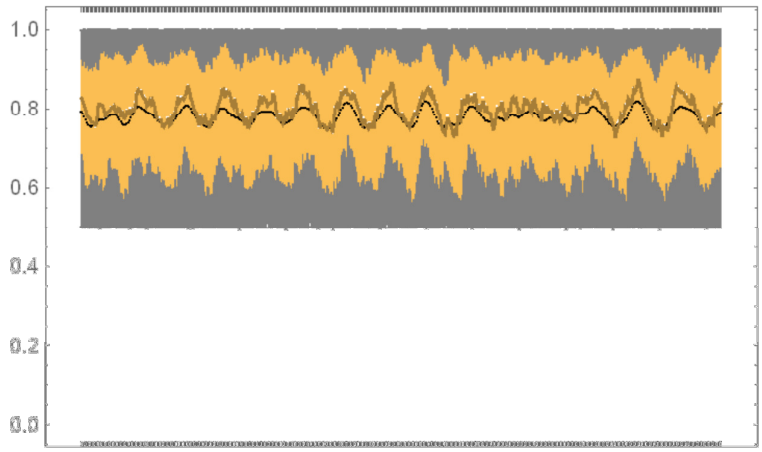

Figure 7 Box-whisker plot of the local order parameter for the J2 staff in the crisis scenario with σ=0.6 but information objects having natural frequencies ω=2. Colour scheme as in Figure 3.

AI Watch-keepers. Now I restore the zero frequency to information objects and remove the noise from the lowest rank watch-keeping staff: OJ2W2, OJ3W2, OJ3S, OJ3WS, and OJ2S. I re-compute solutions to the model at σ=0.6. The local parameters from this are shown, now zoomed in, in Figure 8. Evidently, synchronisation is achieved with all but a small number of instances where some nodes in J2 are caught in a periodic cycle with respect to the main cluster.

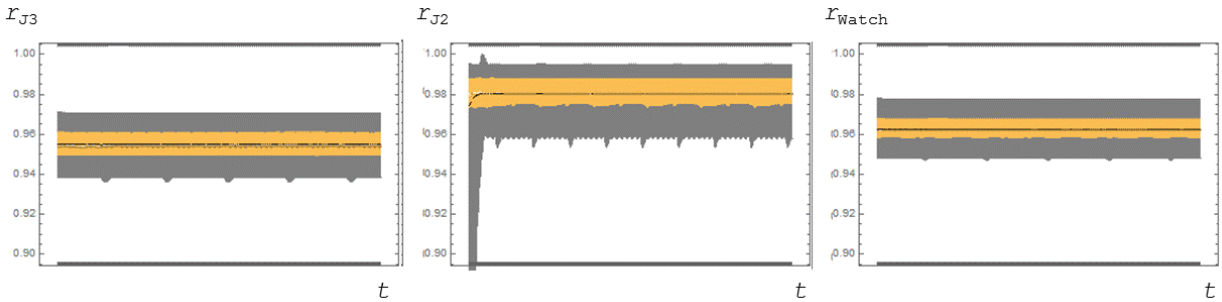

Figure 8 Box-whisker plots for local order parameters Eqs.(3) for crisis scenario at σ=0.6 but with zero frequency information objects and noise removed from low-ranked watch staff.

Successful though this may appear, the instantaneous frequencies of the staff again behave as in Figure 4, converging to below the frequency of the slowest decision-maker, a consequence of the 'dumb' information objects again. It should be self-evident that combining AI agents in place of low-rank staff and smart information objects *but at the pace that human decision-makers can match* is the panacea: synchronisation at the frequency of the fastest human decision-maker. For lack of space I do not present the associated plots from numerical solution.

Adaptive lags. In introducing the formalism of the Kuramoto model I mentioned the scope for including frustration parameters to reflect explicit delays or phase shifts with respect to agents in the network. Here I reintroduce them into the model. Specifically, recent developments with the deterministic system (Brede and Kalloniatis 2016a) have shown that frustrations may be tuned in relation to the (random selected) frequencies so as to improve synchronisation, lower the critical

coupling and drive the system at a collective frequency Ω that (even for the undirected network) *need not equal the mean of the intrinsic frequencies.* This insight is obtained by expanding the interaction $\sin(\beta_j - \beta_i + \varphi_i)$ in Eq.(1) in small fluctuations around a leading behaviour $\beta_i = \Omega t$. This leads (Brede and Kalloniatis 2016a) to a 'recruitment condition' $\Omega = \omega_i - \sigma d_i \sin\varphi_i$ with *d* the degree of node *i*. The sine here may be inverted giving the formula for tuning the frustrations with a preferred value of Ω. Averaging the equation over all nodes for the undirected network shows that $\Omega \neq \overline{\omega}$. A further modification of the model is to allow the frustration parameters to be *dynamical* so as to coevolve with the phases of the Kuramoto model:

$$\begin{aligned} \dot{\beta}_i &= \omega_i + \sigma \sum_{j=1}^{N} A_{ij} \sin(\beta_j - \beta_i + \varphi_i) \\ \dot{\varphi}_i &= \tau \sum_{j=1}^{N} A_{ij} \sin(\beta_j - \beta_i). \end{aligned} \tag{4}$$

One observes that the second equation of (4) vanishes on the right hand side precisely where the phases β are equal and contrastingly the frustrations pick up a maximum rate of change when the phases β are π apart. The parameter τ is a scale that controls the speed of the evolution of the frustrations in relation to the ordinary Kuramoto dynamics. For too small a value of τ, the frustration dynamics is too slow and simply undermines the ordinary Kuramoto synchronisation process (with no frustration); for too large τ, the frustration dynamics is too fast and synchronisation is lost because they dissociate from the Kuramoto phases. Indeed, there is an ideal value of τ at which the frustrations co-evolve with the Kuramoto phases to precisely the tuned values from the above recruitment condition. In this case, however, the value of Ω cannot be completely externally selected but corresponds to one of many equilibria in how the recruitment condition is satisfied as the system undergoes a transient and finds an equilibrium.

In the context of C2, as explained, frustrations represent the staggering of individual decision-loop points that may occur in a well-defined process ordering the sequence of C2 agents completing stages of their decision cycle. Tuning frustrations to individual native frequencies in this context would appear entirely impractical: such quantities cannot be measured for human agents with any precision in advance to enable implementation of a tuning[7]. Dynamical frustrations, however, may represent the presence of *artificial agents* in the C2 system whose role is to stagger, or induce lags, between the human agents such that the total system – now a complex adaptive system that includes AI agents - evolves to a higher synchronised state.

I test now whether such a mechanism can work in the context of 'noisy' humans, in contrast to the theoretical deterministic cases studied in (Brede and Kalloniatis, 2016a). Again, I apply the crisis scenario data with σ=0.6 and the stable Lévy noise. Solving the system numerically for a variety of τ, I find that the point at which frustrations and phases co-evolve to a synchronised equilibrium is at τ=0.7 but where the frustrations have initial conditions drawn randomly in a narrow range in (0,π/4). The result for the local order parameters, again plotting box-whisker analysis for each time step across 100 runs of the model, are shown in Figure 9.

[7] That is why in this paper I draw frequencies from a distribution and solve over an ensemble of instances.

The results are both encouraging and discouraging. The mean values in these plots are barely distinguishable from r=1, very high synchronisation, with narrow quartiles that are barely visible in the plots. However, the exceptional cases are more erratic than for the original case in Figure 5.

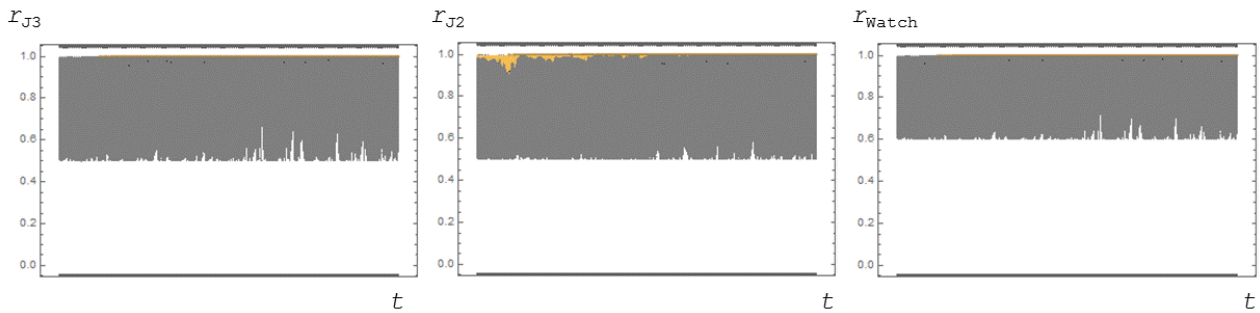

Figure 9 Box-whisker plots for local order parameters Eqs.(3) for crisis scenario at σ=0.6 but with adaptive lags with frustration dynamics constant τ=0.7.

More disturbing is the analysis of the instantaneous frequency for some cases shown in Figure 10. The majority of cases are like the result (a) where the system converges to a frequency close to the mean of the ensemble (information objects again have zero frequency). However some instances show a *negative* frequency, as in (b) – the system is moving backwards! A small number show improved performance in convergence of the human agents to the *fastest frequency*, as in (c). However, in this case there is an undesirable disassociation from the information objects.

The reason for this variety of behaviours lies in the aforementioned multiple equilibria and the range of possible collective frequencies Ω possible. It can be minimised but not eliminated by initializing the frustrations close to zero. Noise plays a role insofar as stochasticity delays the entry of the system into the optimal basin of attraction of an equilibrium (with fastest possible positive Ω), causing the trajectory to reach those corresponding to undesirable collective frequencies. Though at its best this mechanism offers improved behaviour of the C2 system at lower coupling and at faster collective frequency than the mean, the range of possible behaviours renders the system fragile. I return to this problem in the final section.

Seamless mediation of information. There is one more technological scenario that deserves some comment here, where the interactions between human agents no longer require an information artefact but enable seamless transfer of information. I have numerically tested this scenario too using the same parameters as above. The associated (now purely) social networks are provided in the appendix. The results are telling: tuning the coupling, as before, to provide synchronisation for steady-state leads to *improved* synchronisation in the crisis scenario. The collective frequency is now the mean of the human ensemble. However, the coupling to achieve this is *higher*, σ=11, than for the situation with artefact mediated interactions. That information artefacts allow for looser coupling in a system is consistent with Tenet 5 of DSA (Stanton 2006). It is an open question whether the technology that may enable such direct seamless communication between real C2 nodes can also take up the effort from humans in achieving such tight levels of coupling.

9. Conclusions and Discussion

I have proposed a mathematical model for the C2 of a headquarters watch-centre based on the Kuramoto model of network synchronisation of phase oscillators. I calibrated the model using data on how such an organisation conducts steady-state business and reasonable assumptions of other

distributional parameters while testing for sensitivity in these choices. By then applying data for how the staff anticipated handling a crisis I was able to provide event-validation of the model: because of centralisation of the proposed arrangements instabilities were generated in the behaviour, matching the qualitative prediction of Contingency Theory.

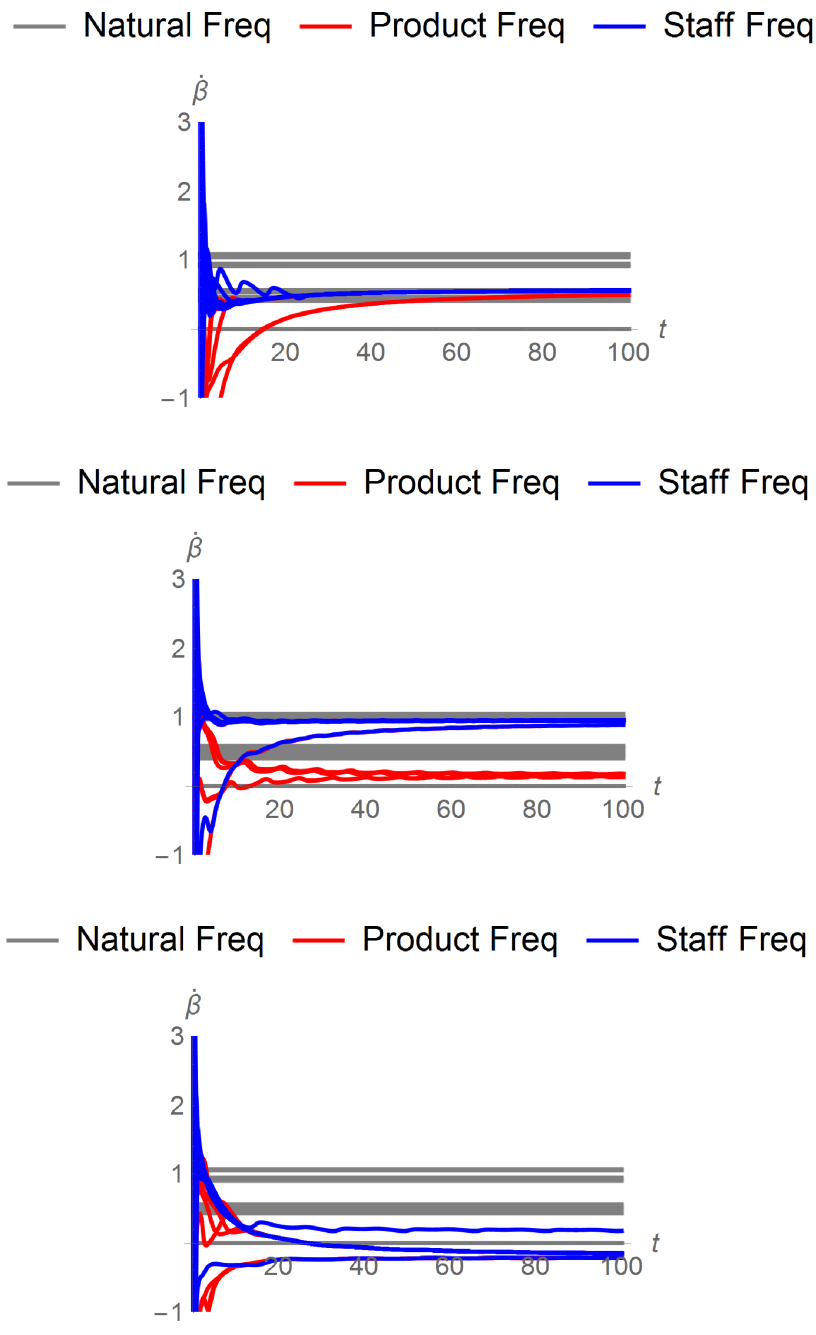

Figure 10 Three instances of the instantaneous frequencies for the crisis scenario data at σ=0.6 with adaptive lags at τ=0.7. Colours are as in Figure 4.

I then explored a range of interventions using AI – 'smart' information objects, AI watch staff and adaptive lags to demonstrate how such a model may anticipate where scenarios for improvement may succeed or enjoy only limited success. Maximum effect was achieved by harnessing AI mechanisms

to one another. In particular, it is not enough to let autonomous components run away with arbitrarily fast decision speeds – they must be paced to the decision speed of the aggregate of those of the human agents to enable coherent *system level* decision making; the mathematical model enables predictive testing of the coherence of such scenarios. That leads to my major conclusion: that a Kuramoto-based model of C2 offers a rich and numerically (and analytically) straightforward means of analysing C2 system performance accounting for the information artefact/synthetic decision-aid mediated connectivity, dynamism and heterogeneity of human-centred Command and Control.

The model's predictive power lies in the ability to change a range of local inputs to the model – network, frequencies – and observe changes of behaviour reflecting the success or failure of real world interventions on a baseline model. Observe that the model amplifies its predictive power by using *distributional characteristics* of human cognitive behavior, encoded in the additive noise; the distributions, naturally call upon a number of parameters that need to be tuned from data, nevertheless they enable a scaling of the system size without corresponding increasing number of input parameters and corresponding reduction of predictive power.

Arguably the point at which the steady-state instance was calibrated may be deemed too high by demanding that *all* instances across 100 runs should frequency synchronise. However, allowing some instances of 'failure' at that point of calibration would have resulted in a coupling rendering the crisis scenario even more fragile. My assumptions for the AI agents may be deemed optimistic, not reflecting the tendency of IT systems to stall or suffer outage. Ultimately, one would wish to measure organisational coupling, namely speed of response by one agent to change in decision state of another, and undertake a higher fidelity validation – such as testing that steady-state exhibits equilibrium at the measured coupling. At the very least, a measure of frequency of use of a given link provides some access to this. The use-case I have undertaken illustrates the type of data that should be measured. Remarkably, for all the acceptance of OODA as a model for C2, cognitive psychology has not yet provided data (to my knowledge). My vision is that this model brings together a range of scientific disciplines, human and mathematical, to enhance the C2 enterprise.

There remains much to be done in the mathematical foundations of the model. The stable Lévy noise used here may be further generalized to 'tempered' Lévy noise (Kullberg and del-Castillo-Negrete 2012) which smoothly cuts-off excessively large jumps allowing a better match to finite statistical data. There is scope for using the Ratcliff model of cognitive diffusion explicitly in this context. The Kuramoto model with co-evolving frustrations is still immature with open questions about how better to control the collective system frequency and rigorous study of the system subject to noise. On this front, however, progress has been made with a different version of the model of Eqs.(4) by which an explicit collective frequency may be selected (Brede and Kalloniatis 2016b). The result is a model that achieves a desired faster collective frequency through internal adaptation to an external control in contrast to an explicitly control driven scenario. Also, the model may be extended to include better representations of the adversary or environment as a network with its own dynamical behaviours and goals, as in the 'Blue-vs-Red' formulation of the model. Finally, through knowledge of the nature of chaos, stability (Kalloniatis and Zuparic 2016) and entropy (Kalloniatis 2014) in such a model one may at last begin to rigorously apply terminology from complexity theory to the domain of C2.

Acknowledgements. I am grateful for constructive discussions with colleagues Yi Yue, Richard Taylor, Mathew Zuparic, Irena Ali, David Cox, Maria Athanassenas and Tim Neville, as well as academics Paul Salmon and Neville Stanton, and my collaborators on the fundamental aspects of this research Dale Roberts and Markus Brede. This work is supported by a three year Chief Defence Scientist Fellowship.

Appendix: Network diagrams

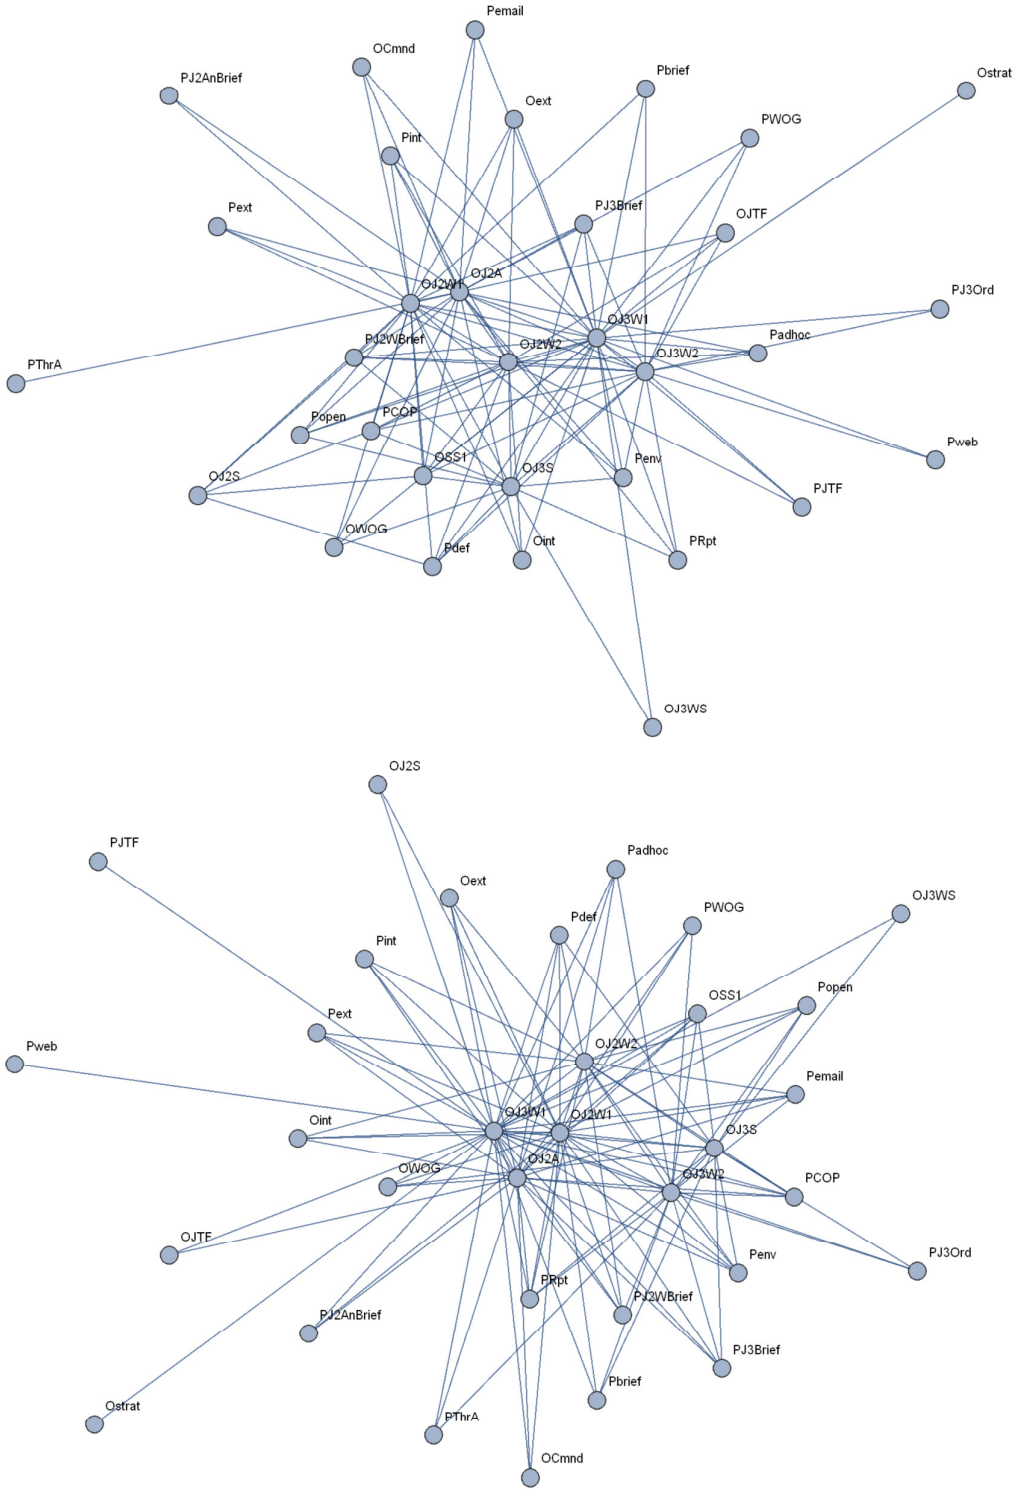

Figure 11 The social-information networks, combining push and pull, for the two scenarios from Ali et al (2014), steady-state (top) and crisis (bottom). These are directed graphs however for clarity here the arrow heads are not shown.

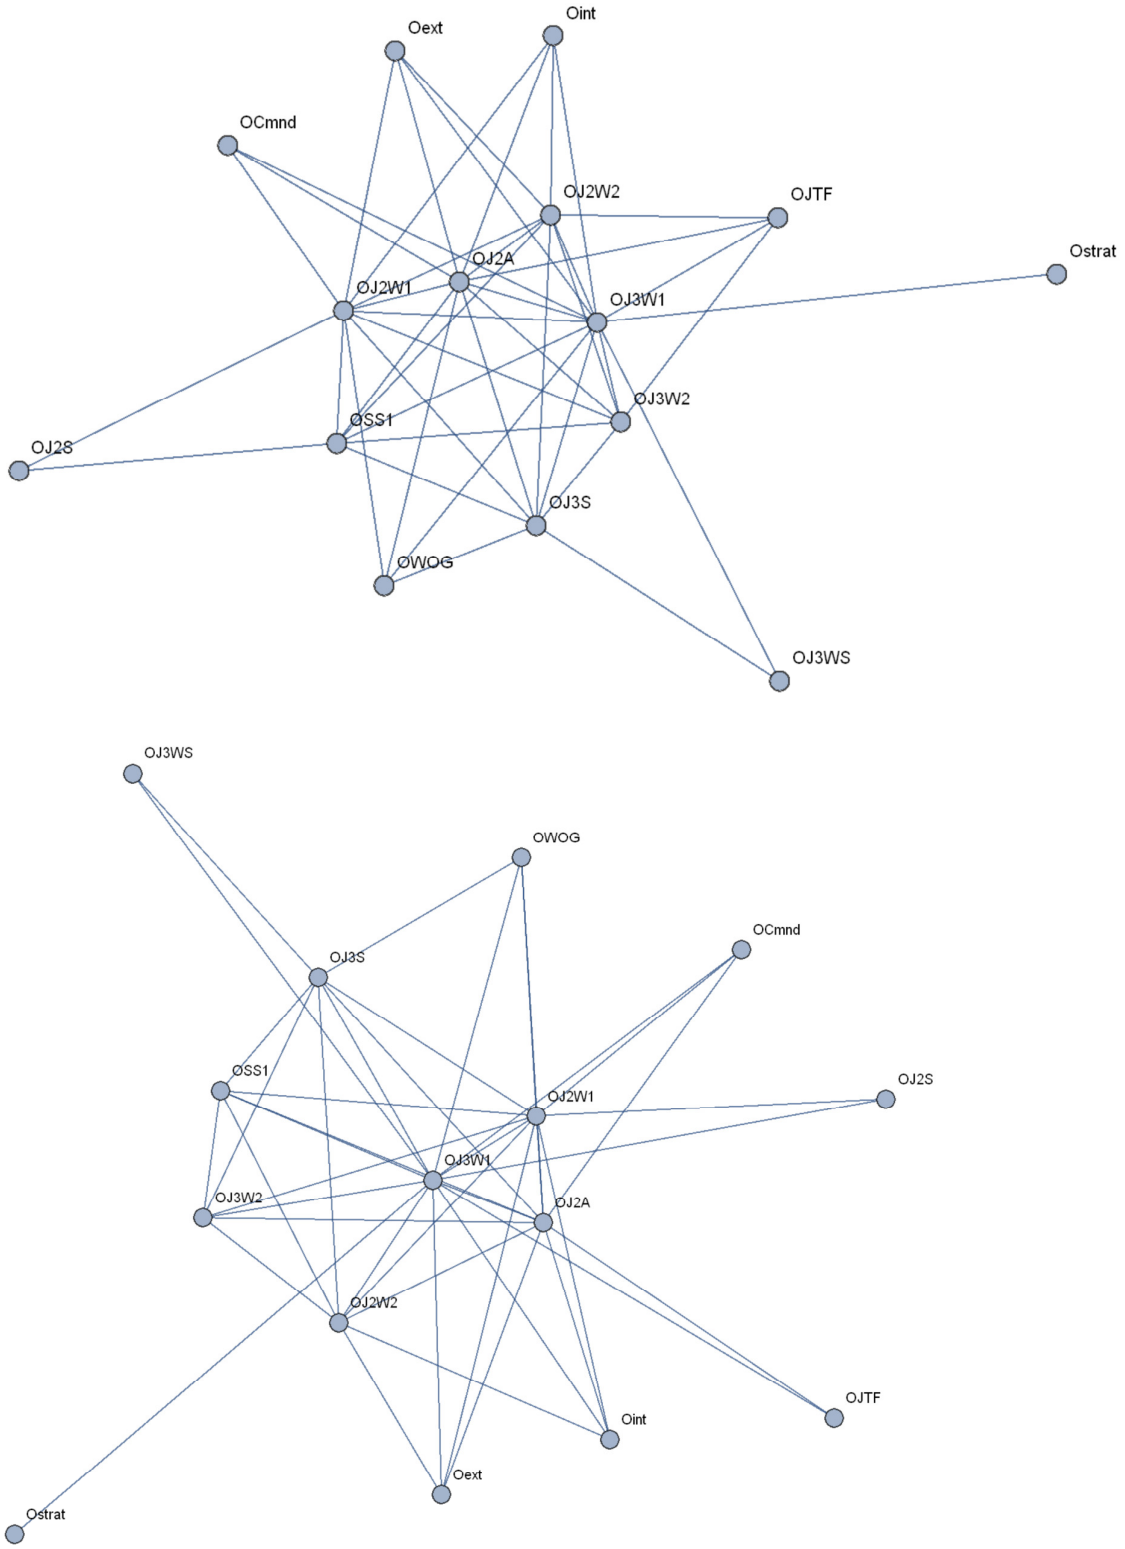

Figure 12 The purely social networks, combining push and pull, and replacing information object mediated links between humans by a single link, for steady-state (top) and crisis scenarios (bottom). These are directed graphs however for clarity here the arrow heads are not shown.

Bibliography

AFIFI, W.A., WEINER, L. ,2002, INFORMATION SEEKING ACROSS CONTEXTS, HUMAN COMMUNICATION RESEARCH, 28 (2), 207–212

ALBERTS, D.S., HAYES, R.E., 2007, PLANNING: COMPLEX ENDEAVOURS, CCRP PUBLICATION SERIES, APRIL, DOD

ALI, I., KALLONIATIS, A.C., KOHN, E., LA, P., MACLEOD, I., NEVILLE, T., ZUPARIC, M., 2014, THE SITUATION AWARENESS WEIGHTED NETWORK (SAWN) MODEL, 19TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM, ALEXANDRIA, VA, USA, JUNE 16-19

BEEKUN, R.I., AND GINN, W.H. 2001A, ORGANIZATION STRUCTURE FROM A LOOSE COUPLING PERSPECTIVE: A MULTIDIMENSIONAL APPROACH, DEC.SCI. 32 (2) 227-250

BEEKUN, R.I., AND GINN, W.H. 2001B, DEVELOPMENT AND TEST OF A CONTINGENCY FRAMEWORK OF COUPLING, J.APPL.BEH.SCI 37 (4) 385-407

BERGER, C.R., CALABRESE, R.J., 1975, SOME EXPLORATION IN INITIAL INTERACTION AND BEYOND: TOWARD A DEVELOPMENTAL THEORY OF COMMUNICATION, HUMAN COMMUNICATION RESEARCH, 1, 99–112

BREDE, M., KALLONIATIS, A.C., 2016 A, FRUSTRATION TUNING AND PERFECT SYNCHRONIZATION IN THE KURAMOTO-SAKAGUCHI MODEL, PHYS. REV. E 93, 062315

BREDE, M., KALLONIATIS, A.C., 2016 B, CODYNAMICS VS CONTROL IN THE KURAMOTO-SAKAGUCHI MODEL, DRAFT PAPER IN PROGRESS

DAWSON, D., REID, K., 1997, FATIGUE, ALCOHOL AND PERFORMANCE IMPAIRMENT, NATURE 388, 235

DE ALWIS, G., MAJID, S., CHAUDHRY, A.S., 2006, TRANSFORMATION IN MANAGERS' INFORMATION SEEKING BEHAVIOUR: A REVIEW OF THE LITERATURE, JOURNAL OF INFORMATION SCIENCE, 32 (4), 362–377

DEKKER, A., 2007, STUDYING ORGANISATIONAL TOPOLOGY WITH SIMPLE COMPUTATIONAL MODELS, JOUR. ARTIFICIAL. SOCIETIES AND SOCIAL SIMULATION, 10 (4) 6 <http://jasss.soc.surrey.ac.uk/10/4/6.html>

DEKKER, A., 2011, ANALYZING C2 STRUCTURES AND SELF-SYNCHRONIZATION WITH SIMPLE COMPUTATIONAL MODELS, PROCEEDINGS OF THE 16TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM, QUEBEC CITY, CANADA, JUNE 21–23

DOERFLER, F., BULLO, F., 2014, SYNCHRONIZATION IN COMPLEX NETWORKS OF PHASE OSCILLATORS, AUTOMATICA, 50 (6), 1539-1564

DONALDSON, L., 2001, THE CONTINGENCY THEORY OF ORGANIZATIONS, SAGE, THOUSAND OAKS

ENDSLEY, M.R., 1995, TOWARD A THEORY OF SITUATION AWARENESS IN DYNAMIC SYSTEMS, HUMAN FACTORS 37, 32-64

ENDSLEY, M.R., 2006, SITUATION AWARENESS, IN: G. SALVENDY, ED. HANDBOOK OF HUMAN FACTORS AND ERGONOMICS, 3RD ED. NEW YORK, WILEY, 528-542

GLASSMAN, R.B., 1973, PERSISTENCE AND LOOSE COUPLING IN LIVING SYSTEMS, SYS.RES. & BEH. SCI 18 (2), 83-98

GRANT, T., KOOTER, B., 2005, COMPARING OODA AND OTHER MODELS AS OPERATIONAL VIEW C2. IN: 10TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM (ICCRTS), MCLEAN, VA, USA, 13-16 JUNE

HARRÉ, R., 1970, THE PRINCIPLES OF SCIENTIFIC THINKING, MACMILLAN, LONDON

HASIK, J., 2013, BEYOND THE BRIEFING: THEORETICAL AND PRACTICAL PROBLEMS IN THE WORKS AND LEGACY OF JOHN BOYD, CONTEMPORARY SECURITY POLICY 34(3), 583-599

KALLONIATIS, A.C., 2008, A NEW PARADIGM FOR DYNAMICAL MODELLING OF NETWORKED C2 PROCESSES, 13TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM, SEATTLE, WA, USA, JUNE

KALLONIATIS, A.C., 2010, FROM INCOHERENCE TO SYNCHRONICITY IN THE NETWORK KURAMOTO MODEL, PHYS,REV.E 82, 066202

KALLONIATIS, A.C., 2012, ON THE 'BOYD-KURAMOTO MODEL': EMERGENCE IN A MATHEMATICAL MODEL FOR ADVERSARY C2 SYSTEMS, 17TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM, FAIRFAX, VA, USA, JUNE 19-21

KALLONIATIS, A.C., ZUPARIC, M.L., 2014, THE J-STAFF SYSTEM, NETWORK SYNCHRONIZATION AND NOISE, 19TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM, ALEXANDRIA, VA, USA, JUNE 16-19
KALLONIATIS, A.C., 2014, ENTROPY AND STABILITY OF PHASE SYNCHRONISATION OF OSCILLATORS ON NETWORKS, ANN.PHYS. 348, 127-143
KALLONIATIS, A.C., ZUPARIC, M.L., 2016, FIXED POINTS AND STABILITY IN THE TWO-NETWORK FRUSTRATED KURAMOTO MODEL, PHYSICA A, 447 21-35
KALLONIATIS, A.C., ROBERTS, D.O., 2016, SYNCHRONISATION OF NETWORKED KURAMOTO OSCILLATORS UNDER STABLE LÉVY NOISE, PHYSICA A, IN PRESS
KHINCHINE A.YA., AND LÉVY, P., 1936, COMPTES RENDUS, 202, 274
KLEIN, G., 1998, 'SOURCES OF POWER: HOW PEOPLE MAKE DECISIONS', MIT PRESS, CAMBRIDGE, MASS, USA
KHOSHBAKHT, H., SHAHBAZI, F., SAMANI, K.A., 2008, J. STAT. MECH: THEOR. EXP.,10, 10020
KIRKLAND, S., SEVERINI, S., 2015, ALPHA-KURAMOTO PARTITIONS: GRAPH PARTITIONS FROM THE FRUSTRATED KURAMOTO MODEL GENERALIZE EQUITABLE PARTITIONS, APPL.ANAL.DISCR.MATH. 9
KULLBERG, A., DEL-CASTILLO-NEGRETE, D., 2012, J.PHYS.A: MATH.THEOR. 45, 255101
KURAMOTO, Y., 1984, CHEMICAL OSCILLATIONS, WAVES AND TURBULENCE, SPRINGER
LLOYD, M., MARKHAM, G., AND DODD, L., 2006, USE OF COMPLEXITY RELATED IDEAS IN UNDERSTANDING MILITARY OPERATIONS AND MILITARY CAPABILITY, 11TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM , CAMBRIDGE, UK, SEPT. 26-28
LÉVY, P., 1937, THEORIE DE L'ADDITION DES VARIABLES ALÉATOIRES, GAUTHIER-VILLARS, PARIS
METZLER, R., KLAFTER, J., 2000, THE RANDOM WALK'S GUIDE TO ANOMALOUS DIFFUSION: A FRACTIONAL DYNAMICS APPROACH, PHYSICS REPORTS, 339, 1-77
MINTZBERG, H., 1979, THE STRUCTURING OF ORGANISATIONS, PRENTICE-HALL, NEW JERSEY, 1979
MIZUMOTO, T., OTSUKA, T., NAKADAI, K., TAKAHASHI, T., KOMATANI, K., OGATA, T., OKUNO, H. G., 2010, HUMAN-ROBOT ENSEMBLE BETWEEN ROBOT THEREMINIST AND HUMAN PERCUSSIONIST USING COUPLED OSCILLATOR MODEL. IN: IEEE/RSJ INT. CONF. ON INTELLIGENT ROBOTS & SYSTEMS. TAIPEI, TAIWAN, PP. 1957 -1963
MOREWEDGE, C.K., KAHNEMAN, D., 2010, TRENDS IN COGNITIVE SCIENCES 14, 435-440
MURPHY, P.J., 2002, FATIGUE MANAGEMENT DURING OPERATIONS : A COMMANDER'S GUIDE, DOCTRINE WING LAND WARFARE DEVELOPMENT CENTRE, VICTORIA AUSTRALIA
NEISSER, U., 1976, COGNITION AND REALITY: PRINCIPLES AND IMPLICATIONS OF COGNITIVE PSYCHOLOGY, SAN FRANCISCO, FREEMAN
OSINGA, F., 2013, 'GETTING' A DISCOURSE ON WINNING AND LOSING: A PRIMER ON BOYD'S 'THEORY OF INTELLECTUAL EVOLUTION', CONTEMPORARY SECURITY POLICY 34(3), 603-624
PALEY, D. ET AL, 2007, OSCILLATOR MODELS AND COLLECTIVE MOTION, IEEE CONTROL SYSTEMS MAGAZINE 27(4), 89-105
PERROW, C., 1984, NORMAL ACCIDENTS: LIVING WITH HIGH RISK TECHNOLOGIES, PRINCETON U.
PIGEAU, R., AND MCCANN, C., 2002, RECONCEPTUALIZING COMMAND AND CONTROL, CANADIAN MILITARY JOURNAL, SPRING, 53-64
PLUCHINO, A. ET AL, 2006, OPINION DYNAMICS AND SYNCHRONIZATION IN A NETWORK OF SCIENTIFIC COLLABORATIONS, PHYSICA A 372(2), 316-325
PROUD, R.W., HART, J.J., MROZINSKI, R.B., 2003, METHODS FOR DETERMINING THE LEVEL OF AUTONOMY TO DESIGN INTO A HUMAN SPACEFLIGHT VEHICLE: A FUNCTION SPECIFIC APPROACH, PROCEEDINGS OF THE PERFORMANCE METRICS FOR INTELLIGENT SYSTEMS (PERMIS), GAITHERSBURG, MD, 1-15
RATCLIFF, R., MCKOON, G., 2008, THE DIFFUSION DECISION MODEL: THEORY AND DATA, NEURAL COMPUTATION, 20, 873-922
RITTEL, H.W.J., WEBBER, M.M., 1973, DILEMMAS IN A GENERAL THEORY OF PLANNING, POLICY SCIENCES 4, 155-169
SAKAGUCHI, H., KURAMOTO, Y., 1986, A SOLUBLE ACTIVE ROTATOR MODEL SHOWING PHASE TRANSITIONS VIA MUTUAL ENTRAINMENT, PROG.THEOR.PHYS.76, 576
SARGENT, R.G., VERIFYING AND VALIDATING SIMULATION MODELS, PROC. OF 1996 WINTER SIMULATION CONFERENCE, PP. 55-64

STANTON, N. ET AL 2006, DISTRIBUTED SITUATION AWARENESS IN DYNAMICS SYSTEMS: THEORETICAL DEVELOPMENT AND APPLICATION OF AN ERGONOMICS METHODOLOGY, ERGONOMICS, 49 (12-13), 1288-311
WEICK, K., 1976, EDUCATIONAL INSTITUTIONS AS LOOSELY COUPLED SYSTEMS, ADM.SCI.QRTLY 21 (1), 1-19
WIENER, N., 1961, CYBERNETICS, 2ND EDITION, MIT PRESS, CAMBRIDGE, MASS
WINFREE, A.T., 1967, J. THEORET.BIOL. 16, 15