

Mobile Health Systems and Emergence



Valerie M. Jones

Volume 8

2015

Number 1

ISSN 1756-2074

About Insights

Insights captures the ideas and work-in-progress of the Fellows of the Institute of Advanced Study at Durham University. Up to twenty distinguished and 'fast-track' Fellows reside at the IAS in any academic year. They are world-class scholars who come to Durham to participate in a variety of events around a core inter-disciplinary theme, which changes from year to year. Each theme inspires a new series of *Insights*, and these are listed in the inside back cover of each issue. These short papers take the form of thought experiments, summaries of research findings, theoretical statements, original reviews, and occasionally more fully worked treatises. Every fellow who visits the IAS is asked to write for this series. The Directors of the IAS – Veronica Strang, Rob Barton, Barbara Graziosi and Martin Ward – also invite submissions from others involved in the themes, events and activities of the IAS. *Insights* is edited for the IAS by Barbara Graziosi. Previous editors of *Insights* were Professor Susan Smith (2006–2009) and Professor Michael O'Neill (2009–2012).

About the Institute of Advanced Study

The Institute of Advanced Study, launched in October 2006 to commemorate Durham University's 175th Anniversary, is a flagship project reaffirming the value of ideas and the public role of universities. The Institute aims to cultivate new thinking on ideas that might change the world, through unconstrained dialogue between the disciplines as well as interaction between scholars, intellectuals and public figures of world standing from a variety of backgrounds and countries. The Durham IAS is one of only a handful of comparable institutions in the world that incorporates the Sciences, Social Sciences, the Arts and the Humanities.

The focal point of the IAS is a programme of work associated with, but not exclusive to, an annual research theme. At the core of this work lies a prestigious Fellowship programme. This programme gathers together scholars, intellectuals and public figures of world standing or world-promise to address topics of major academic or public interest. Their mission is to anticipate the new and re-interpret the old, communicating across and working between disciplinary boundaries.

Every year, the Institute invites as many as twenty highly creative individuals to spend up to three months in Durham. They are located in Cosin's Hall, a magnificent and spacious 18th century mansion which, together with Durham Cathedral and Durham Castle, forms part of Palace Green, dominating the World Heritage Site of Durham Peninsula. During their stay, Fellows engage with departments and colleges, deliver public lectures and seminars, and, above all, join an international community of researchers to address the theme selected for that year. Further details of the IAS and its Fellowship programme can be found at www.durham.ac.uk/ias/fellows


Copyright

The design and contents of *Insights* are subject to copyright. Copyright and Reproduction Rights in all submitted contributions remain with the authors, as described in the Author's Copyright Agreement. Copyright and Reproduction Rights of all other material remain with *Insights*.

Except under the terms of Fair Dealing (UK Copyright, Designs and Patents Act 1988), the user may not modify, copy, reproduce, retransmit or otherwise distribute the site and its contents (whether text, graphics or original research concepts), without express permission in writing from the Institute. Where the above content is directly or indirectly reproduced in an academic context under the terms of Fair Dealing, this must be acknowledged with the appropriate bibliographical citation.

The opinions stated in the *Insights* papers are those of their respective authors and do not necessarily reflect the opinions of the Institute of Advanced Study, Durham University, or the staff and students thereof.

MOBILE HEALTH SYSTEMS AND EMERGENCE

Changes in the age distribution of the population and increased prevalence of chronic illnesses, together with a shortage of health professionals and other resources, will increasingly challenge the ability of national healthcare systems to meet rising demand for services. Large-scale use of eHealth and mHealth services enabled by advances in ICT are frequently cited as providing part of the solution to this crisis in future provision. As part of this picture, self-monitoring and remote monitoring of patients, for example by means of smartphone apps and body-worn sensors, is on the way to becoming mainstream. In future, each individual's personal health system may be able to access a large number of devices, including sensors embedded in the environment as well as in-body smart medical implants, in order to provide (semi-)autonomous health-related services to the user. This article presents some examples of mHealth systems based on emerging technologies, including body area networks (BANs), wireless and mobile technologies, miniature body-worn sensors and distributed decision support. Applications are described in the areas of management of chronic illnesses and management of (large-scale) emergency situations. In the latter setting BANs form part of an advanced ICT system proposed for future major incident management; including BANs for monitoring casualties and emergency services personnel during first response. Some challenges and possibilities arising from current and future emerging mHealth technologies, and the question of how emergence theory might have a bearing on understanding these challenges, is discussed here. 

Introduction

Mobile health (mHealth) systems are electronic health (eHealth) systems which utilise mobile technologies to provide mobile health-related services to users. At the University of Twente we apply the emergent technologies of body area networks (BANs), wearable sensors and wireless communications in the healthcare domain in order to provide mHealth services to patients and professionals.

Body Area Networks

We define a BAN as a 'network of communicating devices worn on, in or around the body providing mobile services to the wearer', thus including the possibility of implanted devices as well as body-worn devices and devices carried by users (e.g. in a pocket) or worn as part of 'body furniture' (e.g. jewellery or spectacles). A mobile platform such as a personal digital assistant (PDA) or smartphone handles storage, processing and communications. This generic concept of BAN can be specialised for particular application domains; we focus on BANs to support healthcare, referring to this class as health BANs. Health BANs can provide local mobile health services such as patient monitoring, (bio)feedback, coaching and even treatment. These BANs may include medical devices such as biosensors as well as general purpose devices (e.g. alarm buttons). A health BAN then might consist of a smartphone and a set of body-worn sensors which measure the patient's biosignals (blood pressure, heart rate and blood glucose, for example). Sensor data may be processed locally on the BAN or sent to a remote system for processing, or a combination of the two.

Where communication with a remote system is involved, for instance transmission of biosignals to a hospital or to a (mobile) professional, the mHealth services can be designated *teleservices* (examples are telemedicine services such as telemonitoring and teletreatment). An early example we proposed of teletreatment mediated by a BAN was the possibility of applying a feedback-control loop so that a medication delivery device (e.g. implanted insulin pump) is controlled on the basis of ongoing monitoring of biosignals (e.g. blood glucose) (Jones et al., 2001). Several examples of applications of health BANs are described below, in order to illustrate the generic concepts and the range of possibilities for mHealth applications.

Body Area Networks in Chronic Care

From 2002 to 2004 during the IST project MobiHealth (<http://www.mobihealth.org>) the health BAN concept was specialised for different clinical applications, and a number of health BANs prototyped and trialled with patients suffering from chronic conditions including cardiac arrhythmias, chronic obstructive pulmonary disease (COPD) and mental health problems. Additionally, a pregnancy monitoring BAN was developed; this example is used to illustrate the principles. The pregnancy monitoring trial was conducted with the collaboration of the gynaecology department at a hospital in the Netherlands, Medisch Spectrum Twente. The target group was women with high-risk pregnancies who would normally require in-hospital monitoring for a period of three days at a time. The pregnancy BAN had the following hardware configuration: the mobile platform was a PDA (IPAQ) and the BAN devices comprised five electrodes for measuring electromyography (EMG) and an alarm button. The objective was to detect premature labour or foetal distress from the EMG signals. For reasons of patient safety the trial group consisted of patients with normal pregnancies. The BAN was fitted by the gynaecologist and the patient was then able to be monitored from home or any other location, whilst following normal daily life activities, with biosignals and any alarm indications transmitted in real time from the BAN to the care team at the hospital over wireless cell phone communications (general packet radio service (GPRS) and Universal Mobile Telecommunications System (UMTS)).

This example illustrates the general concept of mobile monitoring by means of a health BAN; of course the exact hardware configuration and software application has to be designed and implemented based on the specific requirements of each clinical application. Variants of the BAN, using different sensor sets, were prototyped during the MobiHealth project, with nine different specialisations of the health BAN trialled on patient groups in four European countries.

MobiHealth illustrated the feasibility and utility of transmitting biosignals wirelessly over 2.5 and 3G. In later Dutch and European projects more advanced BAN applications for chronic disease management were developed, introducing teletreatment alongside telemonitoring and involving more and more sophisticated processing, analysis and interpretation of biosignals in the light of general and patient-specific clinical data and knowledge as well as various kinds of context data.

Further development of BANs for chronic disease management is currently conducted in the European project MobiGuide (<http://www.mobiguide-project.eu>). The MobiGuide patient guidance system supports the patient and the medical team caring for them in adhering to best evidence as encapsulated in clinical guidelines. Moreover, it supports communication between them, information sharing and shared decision-making between patient and clinician. The MobiGuide BAN is part of a distributed decision support system (DSS). The knowledge base of the distributed DSS is based on the knowledge encapsulated in clinical guidelines.

Clinical guidelines are documents which bring together the best and latest scientifically-proven knowledge about how to manage and treat a particular condition and as such represent current medical consensus. They are developed by panels of medical experts who review evidence from clinical trials and scientific literature in order to define best practice and support evidence-based medicine.

Two patient groups are the especial focus in MobiGuide: pregnant women who develop diabetes during pregnancy (gestational diabetes mellitus or GDM) and patients suffering from a cardiac arrhythmia (atrial fibrillation or AF). The AF application allows patients to manage their condition from home or while on the move under the guidance of clinical knowledge adapted to their individual treatment plan and situation. AF patients can record incidence and severity of symptoms on their smartphone and receive appropriate guideline-based but personalised advice from a distributed decision support system. The AF guideline, which is written in English, has been formalised to produce a computer interpretable version (a computer interpretable guideline or CIG). This CIG is customised to include social and technological context and forms the major part of the knowledge base in the decision support system. This CIG is then personalised for each individual patient by the cardiologist working together with the patient to decide which options suit their individual health condition and situation.

Measurements such as blood pressure, heart rate and ECG are uploaded automatically by the sensors of the MobiGuide BAN. Other measurements, such as weight, blood pressure, international normalised ratio (INR) values, dietary information and exercise can be entered manually by patients who can view their own data in a personal logbook running on their smartphone. AF patients can also use the system to monitor exercise sessions and receive warnings and advice, for instance when they exceed their personal recommended limits.

The GDM application allows the patient to enter her daily measurements and receive guidance from the system, which helps in her daily management of her blood glucose levels by means of diet, exercise and, in some cases, insulin therapy, according to the GDM guideline and her current treatment regimen and context. Functions include manual entry of insulin taken and ketone levels and manual or automatic entry of blood glucose and blood pressure readings, depending on whether the patient uses Bluetooth-enabled blood pressure and blood glucose sensors or not. MobiGuide is quality aware; when the patient enters an out-of-range value, for example, the system detects the error and displays a message querying the measurement and offering a second try at entering the value. As with the AF case, the patient can view their data in their personal logbook.

Behind these end-user applications there is a complex distributed system incorporating a distributed knowledge-based system. The BAN components comprise the mobile part of this distributed system. The MobiGuide BAN is based on the Samsung Galaxy S4 smartphone. The following sensors are used: the BioHarness 3 sensor belt from Zephyr, the Omron 708-BT blood pressure monitor, a glucometer for measuring blood glucose, plus two specially developed software detectors; one for detecting atrial fibrillation and another for measuring physical activity. When suitable sensors that can transfer data automatically become available and affordable, they can be integrated into the MobiGuide system, removing the need for manual entry and making the system easier for patients to use and less error prone.

Body Area Networks in Emergency Care and Emergency Management

During the first health BAN project, MobiHealth, the team also developed and trialled two BANs for use in an acute setting: a trauma patient BAN, to be applied to casualties by ambulance paramedics for monitoring casualties in medical emergencies, and a BAN for health professionals, to be used by the ambulance paramedics attending the casualty. The trauma BANs, first proposed in Jones et al., 2001, were developed and trialled in collaboration with the Dutch ambulance service and traumatologists at a Dutch regional trauma centre, Medisch Spectrum Twente. Two BANs were prototyped and trialled: a trauma patient BAN and a BAN for paramedics. The trauma patient BAN included 4-channel ECG, a respiration sensor and a finger clip pulse oximeter. Other measurements (fluids, blood pressure and pupil reaction) were entered manually. The paramedic BAN enabled audio communications and transmission of images from the scene. The original intention was to include real-time video communications but in MobiHealth only still images were transmitted.

This emergency scenario was scaled up during 2004–6 in the IST project MOSAIC. The remit of MOSAIC, and the associated Ambient Intelligence At Work initiative of the EU, was not to develop technology but, firstly, to create futuristic visions of applications of future and emerging technologies which would enable ambient intelligence in the workplace; and, secondly, to develop roadmaps for future research and development towards implementation of these envisioned future applications. In this context a number of visions involving application of ambient intelligence was developed and analysed; one of these visions related to emergency management during the first response to a major incident. The vision was illustrated by a scenario – the MOSAIC Major Incident Scenario – together with an analysis of the future and emerging technologies that vision implied (Jones and Saranummi, 2010).

The MOSAIC Major Incident Scenario envisaged the large-scale use of advanced versions of the MobiHealth BANs in future emergency settings, namely the trauma patient BAN to support triage and monitoring of casualties, as well as use of professional BANS (future versions of the MobiHealth paramedic BAN) by emergency services personnel to support communications, tracking and health and wellbeing monitoring for personnel at the scene and to support telepresence and augmented reality experience of the scene for control-centre personnel up the command chain (Jones et al., 2005). In addition, we envisaged opportunistic construction of emergency communications networks where the BANS of the emergency services personnel and the vehicular networks of the emergency services vehicles could link together by ad hoc networking to construct an emergency communications network at the scene to plug gaps in local communications infrastructure damage (Jones et al., 2011).

The first response phase in major incident management was further researched, in collaboration with various partners in Durham, during a three-month IAS research fellowship in 2014. In particular, the state of the art of current ICT systems used by the fire service and police in Durham was investigated, and a preliminary analysis of requirements on specialised BANs for police and firefighters was conducted. This work complements the previous research conducted in collaboration with trauma surgeons, paramedics and hospital emergency departments during past projects on emergency management which focused on medical services.

Major Incident Management – Current Situation in the UK

In the first stage of a major emergency such as a natural disaster, major transportation accident or major terror attack, multiple casualties as well as massive damage to infrastructure may be involved and a coordinated response by the emergency services and other first responders is required.

First response refers to the phase where action is taken to address the immediate problems faced in an emergency situation, including rescue and treatment of casualties and managing the situation on the ground. Well-managed first response depends on prior phases involving risk analysis, planning and resilience and capacity building, and is followed by many other phases from ensuring business continuity through to reconstruction. Here we focus only on the first response phase.

During first response, each of the emergency services might require involvement by multiple teams from different regions (service 'intraoperability'). Further, involvement of the different services (police, fire service, ambulance and possibly coastguard) and other first responders requires cross-service cooperation (service 'interoperability') (https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/5904/nationalframework.pdf) at operational, tactical, strategic levels. In some cases major incident management requires an international response and hence cross-border cooperation and coordination of first responders.

Analyses of major incidents from around the world (e.g. The 9/11 Commission Report; Zoraster, 2006; Report of the 7 July Review Committee, 2007; Wood-Heath and Annis, 2004) show that communication and coordination and service intra- and interoperability continue to be persistent challenges in major incident and disaster management. Similar challenges are faced in the UK as demonstrated by the findings of the Pollock Review (Pollock, 2013) which also concludes that change is needed in the whole socio-technical context of disaster management.

The UK government's Joint Emergency Services Interoperability Programme (JESIP) is a response at national level to addressing the problem of service interoperability in the context of the provisions of the Civil Contingencies Act 2004 (http://www.legislation.gov.uk/ukpga/2004/36/pdfs/ukpga_20040036_en.pdf). From the ICT perspective, organisational intra- and interoperability need to be supported by compatible and interoperable ICT systems which moreover are resilient in the face of infrastructure damage and possible cyber attack in rapidly evolving emergency situations.

In the UK and elsewhere the ICT systems, sensor networks and underlying network infrastructures have themselves become critical infrastructures supporting logistics, situation assessment, communications and information sharing in emergency management as well as in everyday life. However, these systems themselves have vulnerabilities and may malfunction, suffer damage or themselves be targets for (cyber) attack. As a result, in the very challenging and often rapidly evolving emergency situations, difficulties relating to communication and coordination may be exacerbated by accidental or deliberate disruption of communications.

In fact, the design, development and operation of the ICT systems so necessary to support command and control, situation assessment, communication and coordination and information sharing are themselves the source of many technical and non-technical challenges, some of which are general to (complex) ICT systems and some of which are especially relevant for safety critical systems in general and therefore to mHealth/eHealth systems and to systems supporting

emergency management. As with organisational interoperability, interoperability of ICT systems continues to be one of the major ICT-related technical challenges.

Some ICT Challenges and Issues Arising

There are many technical and socio-technical challenges arising out of recent developments in mHealth systems and the prospect of mass roll-out of mHealth services in emergency and routine situations. They include issues relating to:

- Correctness, safety and security: assurance of quality of software, systems and overall *good behaviour* (a prerequisite for patient and public safety) of distributed systems which increasingly rely on a plethora of autonomous intelligent components. Safety/security issues are even more critical in *emergency settings* and where *teletreatment* as well as *telemonitoring* is involved.
- Reliability and dependability: resilience, robustness.
- Interoperability, scalability, flexibility.
- Usability: natural interfaces (systems must assist without interfering with task(s) in hand); mobility (not tethered to vehicles), handsfree.
- Ethical issues surrounding privacy, routine surveillance; public and private decisions regarding trade-offs between surveillance and control versus privacy and autonomy.
- Determination of policy, and change management, in relation to: healthcare organisation and delivery, professional culture and public acceptance, in the face of (r)evolution in the delivery of care enabled by advances in technology and driven by stakeholder interests including economic, demographic and political pressures.

In the following section we concentrate on the first issue, namely of quality assurance of complex ICT systems.

Emergent Behaviour in Complex Distributed (ICT) Systems

In the foregoing, two connections with the themes of emergence have been alluded to: firstly, the use of emergent technologies (including BANs, wireless communications, miniature sensors, clinical decision support systems) to provide mobile health services to patients in chronic and acute settings; and, secondly, the use of these and other emergent technologies to support emergency management during first response to major incidents. At this point we introduce a third connection, namely the issue of emergent behaviour of complex (ICT) systems, which has implications relating to correctness and hence to reliability, dependability and safety of ICT systems including mHealth systems in chronic care and in emergency management. The challenges of quality assurance mentioned in the previous section are complicated by the possibility of complex ICT systems exhibiting emergent behaviour.

Some ICT systems are in fact designed, and intended to produce (or give the appearance of producing), some kind of emergent behaviour. For example, non-deterministic systems may be designed to produce different outcomes in the same situation (e.g. for reasons of fairness in a multi-user operating system) and neural networks and other AI systems may utilise heuristics or learning algorithms with the production of 'emergent' behaviour as the overt intention. Some such systems are provided with underlying, sometimes non-algorithmic, mechanisms which enable discovery, recognition or synthesis or production of (new) patterns or phenomena when supplied with sufficient and relevant sets of data and examples and some kind of mechanism or

engine for driving the process. In these cases 'emergent' behaviour is intended, expected and will usually be welcomed. It can be argued, however, that such systems are all deterministic at a certain level and hence in principle can be explained by reduction, providing we have the means to access all the relevant information at all relevant levels of abstraction.

Whilst emergent behaviour and unpredicted outcomes are the intention in some cases, in many other cases, however, the possibility of encountering unexpected behaviour from an ICT system is unintended and undesirable. Though the designed or intended 'emergent' behaviour is interesting and deserves study in the context of emergence theory, it is the unintended, undesirable 'emergent' behaviour of complex ICT systems that is the focus here. When managing a health condition or an emergency situation, for instance, one thing we do not need is unexpected, unintended emergent behaviour from the systems we have designed. This applies also to any serious ICT application, serious in the sense that undesirable consequences and negative impact may follow from malfunction. mHealth systems, like all safety-critical or life-critical systems, are serious applications in this sense. They are also (often) complex ICT systems. Furthermore, they are likely to become more widespread, more complex and more autonomous with time. It is obvious that for reasons of patient and public safety, we need to understand how to manage these phenomena in order to build resilient and trustable systems.

Why is it that software systems, which after all are artefacts of our own creation, which we have carefully designed using formal mathematically-based notations, implemented on machines of our own making and tested rigorously, can give us nasty surprises by suddenly and apparently at random exhibiting unwanted and unexpected behaviour?

The fact is that software systems are some of the most complex artefacts created by humans. It is not possible to create a bug-free (non-trivial) program and it is not possible to prove (in the general case) that a program is correct. The potential for nasty surprises is very common because 'bugs' (errors) are ubiquitous in software: the industry average is claimed to be 'about 15–50 errors per 1000 lines of delivered code' (McConnell, 2004) and this holds irrespective of which programming language is used.

The question of unintended emergent behaviour of software systems is a well-known problem in computer science, and creates a major issue for software (and systems) engineering, namely the theoretical problem of verification of correct behaviour of software (and, by implication, verification of correct behaviour of (complex) ICT systems). The 'inconvenient truth' is that even standalone, deterministic, algorithmic systems can misbehave.

We would always like to be able to answer the question 'Does my (computer) program do what it should do (and never what it shouldn't do)?' in the affirmative. Unfortunately we cannot. To rephrase and generalise, the answer to the question 'Is there a decision procedure that can tell me, for any program, if the program is correct?' is 'No', and this assertion is provably correct. It is entailed amongst others by the well-known 'halting problem', which can be formulated informally as: *it is provable that there is no decision procedure which can tell me, for any program, that the program will terminate for every possible input.* The 'Halting problem', also known as the Church-Turing thesis, is one of many provably unsolvable problems in computer science, in this case the problem of undecidability ('*Entscheidungsproblem*' or 'decision problem'). The proof was provided in 1936 independently by Church using lambda calculus and by Turing using the concept of the Turing Machine (Turing, 1937). Although the name 'halting problem' came later, the problem itself is older and its origins have been traced back to Leibniz.

For our purposes, the point to note is that even simple properties of programs such as termination cannot be proven in the general case. Although methods for proving correctness have been developed (in formal verification mathematics and formal logic are used to prove certain *pre-stated desirable properties* of a program), and can be successfully applied in specific cases, these methods are in many cases not practically applicable. Further, there is no guarantee that they will find *unknown or unexpected* problems. Furthermore, in the general case complete proofs of correctness will always be provably impossible, irrespective of how powerful our computers (used in this case as proof engines) may become in future. The reality is that most real programs cannot be formally verified in a reasonable time frame, and some specific programs can provably never be formally verified (for example by exhaustive testing or by proofs of correctness) for a number of reasons, one of which is that in some cases the search space is infinite. But the Church-Turing thesis tells us that in the general case there can be no oracle on program correctness, because there can be no oracle on program termination. These observations lead us to note a distinction between the nature of the propositions. The statement above: 'methods for proving correctness ... are in some cases not practically applicable' relates to a contingent truth. The implication is that in future, given sufficient time and increases in computing power, individual members of this class of proof problem could change from being intractable to being tractable. The halting problem on the other hand appears to be a universal truth (it can be proved in principle that no decision procedure which acts as a generic oracle on program termination can exist, because it entails a logical contradiction). This can never be solved by throwing more computing power at the problem. Instead *software engineering relies on best effort attempts at attaining higher levels of confidence in the quality of developed software through application of systematic development methods including rigorous testing.*

Even standalone programs (programs running on a single computer) do not execute in a vacuum; they run on top of many layers of system software which mediates between the high level software and the lowest levels of code, which interact directly with the hardware which executes the program. The complexity of overall system behaviour increases further and by orders of magnitude when we consider distributed systems. In a distributed system a program runs in parallel across multiple, possibly remote, nodes of a computer network. In this case the (distributed) system is also subject to the vagaries of the networking environment, including failure of links, routing protocols which send data by different routes in an apparently ad hoc manner (actually it is deterministic at some level) and variable transmission delays which together introduce a new class of problems relating to timing: 'concurrency' problems. Behaviour of networks becomes predictable only in probabilistic terms; indeed vendors' technical specifications and service providers' service level agreements (SLAs) for performance of networks and computer systems are expressed in statistical measures relating to failure rates, such as mean time between failures (MTBF), availability (the proportion of time a system is expected to be operational) and reliability (the probability that a system will produce correct outputs within some given time frame).

The sheer complexity of today's networks in terms of numbers of nodes and hence potential numbers of connections between nodes, plus the increasing numbers and diversity of devices which are directly or indirectly connected, only exacerbates the problem. Systems which test valid in the lab become embroiled in complex webs of interconnected devices when released into the wild and system failure is often the result.

A report commissioned by ENISA (European Union Agency for Network and Information Security) found, amongst others, that database interoperability and device interoperability are key priorities for research on current and emerging network technologies. The report comments

on sensor networks, key users of which are police, immigration, security services and emergency services, as follows:

Many sensor network systems are procured specifically for a given purpose and devices are tested systematically to verify that the real effects of their operation in the system comply with the stated purpose and other requirements of the purchaser. This verification applies up to the instant at which the system is turned over to the customer; and even then, what works in a development laboratory or test site may not continue to work when a customer installs personal devices alongside the system's devices. After that, the system may evolve in many ways and one outcome is that additional or replacement devices compromise its function. This problem is a major headache for appliance manufacturers in home and building systems and in healthcare (Gorniak, 2010).

Note particularly the observation that once the system is in use and exposed to a complex networking environment 'the system *may evolve in many ways* and one outcome is that additional or replacement devices *compromise its function*' [my italics]. The report contains a detailed analysis of this and other risks and vulnerabilities with guidelines for design choices to mitigate the effects. In general, in fact, *mitigation is the best that we can expect in terms of remedy for the behavioural problems of complex ICT systems*.

Computer scientists and engineers have developed analysis and design and development methods for aiming for the highest (software) quality possible in the circumstances. These methods include systematic design and development methodologies and tools, (partial) verification approaches such as formal testing, work on quality of data and quality of service and work on reliability and dependability of systems. Industry adopts (usually with a lag time of the order of years) (some of these) approaches. But the fact is that no method can guarantee correct operation in all circumstances, so strategies for managing and ring-fencing possible problems that may arise need to be included at the design stage.

Given that (unintended) emergent behaviour of complex systems is a major problem for mHealth systems in general and the emergency setting in particular, can we get some purchase on this problem, or at least some further insight, by examining the issues in terms of emergence theory and consideration of which kind of emergence we are dealing with in relation to complex behaviour of ICT systems?

Some Observations on Emergence

Disasters and emergencies, whether they originate in and/or impact on physical reality or cyber space (or both), often occur suddenly and apparently without warning, in other words they seem to emerge unexpectedly and from nowhere. Herein lies one connection between emergence and emergency. When phenomena are unplanned and unexpected and (for now) appear to be inexplicable (and especially if they are disruptive) there is a temptation to see them as emergent. One explanation in some cases is given by catastrophe theory, which describes a class of events which are caused by gradual incremental change in some variable or set of variables which, when a critical value or state is reached, may trigger some catastrophic event. Imagine an off-road vehicle ascending a mountain side, where the inclination of the slope is increasing gradually with altitude. As the angle of inclination continues to increase, at a certain point the vehicle will suddenly overbalance and fall backwards into the abyss below, because the point has been reached where the position of the centre of gravity of the vehicle relative to the ground supporting it has passed beyond a critical value. The increase in the angle of inclination was continuous, no quantum leaps were involved in the changing value of that

variable, but a tipping point, literally and metaphorically, has been reached, with sudden and dramatic consequences. (You may wish to try this at home, but only with a Dinky toy.)

Anderson describes two different ways of thinking about emergency: 'the sudden irruption that emerges from within life without warning [...] and the incubating disaster that is already ongoing beneath awareness. [...] Terrorism being the paradigmatic example of the first type of event, and climate change being the paradigmatic example of the second' (Anderson, 2012).

Is the fall of the off-roader a sudden irruption or an incubating disaster? It appears at first sight to be both at once. On the one hand, there is no warning to the driver of the imminent fall, on the other hand, the physical principles involved are well understood, and the event could have been predicted on the basis of theory together with the pertinent data. We need to consider this question: 'sudden irruption or incubating disaster?' in relation to emergencies of all kinds, if only because if emergencies can be *predicted* we can prepare for them or perhaps avert them. If we look at the notions of weak and strong emergence of Chalmers, we may get more leverage on this question. Chalmers distinguishes two different concepts of emergence: *strong emergence* and *weak emergence*, characterising strong emergence thus:

We can say that a high-level phenomenon is *strongly emergent* with respect to a low-level domain when the high-level phenomenon arises from the low-level domain, but truths concerning that phenomenon are not *deducible* even in principle from truths in the low-level domain (2006).

and associating this interpretation with philosophical debate on emergence dating from the 1920s. In contrast he identifies the concept of weak emergence with more recent debate from the scientific community, characterising it thus:

We can say that a high-level phenomenon is *weakly emergent* with respect to a low-level domain when the high-level phenomenon arises from the low-level domain, but truths concerning that phenomenon are *unexpected* given the principles governing the low-level domain (2006).

One consequence of this distinction between strong versus weak emergence seems to be that strongly emergent phenomena cannot ever be predicted from knowledge of the lower levels because there can be no complete reductionist explanation. For weakly emergent phenomena explanation (and therefore prediction) may be difficult, but the possibility exists; this then becomes a question of knowledge and *knowability*. In other words the predictability problem reduces to an epistemological question. Do we have an adequate understanding or model of the underlying processes? Do we have (or can we access) all the relevant data? If the answer to either of these questions is *No*, then prediction is not possible, now, for pragmatic reasons but may become possible in the future. If the answer to both of these questions is *Yes*, then prediction is possible, if we can afford it and if there is sufficient time. Yesterday we worshipped the mountain, but today we finally understood volcanology, so tomorrow we will use seismography to detect and predict eruptions. This contrast between unpredictability and possibility of prediction (though possibly intractable at present) is reminiscent of the contrast between provably unsolvable problems versus the merely intractable problems (intractable at a given point in time) of computer science.

Chalmers poses the question, 'Are there any strongly emergent phenomena?'. He concludes that there are, but can find only one candidate: consciousness. In a small way I searched for examples of strongly emergent phenomena, starting by looking for classificatory dimension such as emergent behaviour in natural systems (such as insect colonies, weather events, natural disasters, virology, epidemiology), emergent behaviour in artificial systems (e.g. artificial life), human and social phenomena (e.g. man-made disasters, crime, terrorism), natural language (emergence of spoken language as a natural and ubiquitous process in contrast to the human

artefact of written language) and in artificial or formal languages (e.g. computer programming languages, language of mathematics).

Some possible dimensions for consideration were identified: whether the phenomena are natural or man-made (evolved or designed); intended or unintended (planned or unplanned); predictable or not (expected or unexpected); the degree of predictability (probabilistic and fuzzy systems, uncertainty, non-determinism); and whether they are benign or malign (desirable or undesirable). The last refers to the observation that perhaps we pay more attention to those events which rate higher according to some function of the seriousness, to us, of their impact on things which concern us. Applying these concepts I found several phenomena that on examination seemed to be weakly emergent. It is now known that swarm behaviour, as exhibited by flocking birds and some species of fish shoals for example, can be described in some cases by very simple rule sets, such that computer programs can easily be written which emulate (at any rate the appearance of) the complex behaviour patterns. Regarding predictability of emergencies, we could argue that if we knew all about all the underlying mechanisms and contributing factors (including the human psychosocial elements) (that is, if we had a kind of omniscient Health and Safety Executive (HSE)) emergencies would be predictable. Indeed post-hoc analyses of air accidents, fires, road accidents, nuclear accidents (e.g. Three Mile Island) and fatal accidents involving medical equipment (e.g. Therac 25) seem to provide compelling and apparently comprehensive analyses of the contributory factors and evolution of complex emergencies; this seems to suggest that in principle these incidents could have been predicted.

Natural disasters and extreme weather events are sometimes predictable, but not always reliably. Some of our models have strong predictive power (especially in the short term) but performance tends to decrease when we peer further into the future. Models are usually simplifications, often probabilistic, and 'noise' or error can overwhelm 'signal' the further ahead we look. As chaos theory shows, small initial effects can be amplified with time and space and move from having negligible to overwhelming impact on accuracy and reliability of results.

Finally, is emergent behaviour in complex ICT systems a case of weak or strong emergence? And what are the consequences for design of ICT systems, and planning for and managing ICT disasters?

There is no theoretical limit on the complexity of the computational models (or computer programs) that we can create (though in real life there will be practical constraints contingent on the particularities of any real computing environment). This holds not only for the informal everyday understanding of complexity, but also in terms of the formal mathematical definitions of computational complexity and network complexity. Might emergent behaviour of complex systems in general, and complex ICT systems in particular, qualify as a candidate for strong emergence? Or do we simply face everyday weak emergence problems with ICT systems, some of which may seem intractable, but with effort and technological advances may become tractable? This hope fades somewhat when we remind ourselves that the theoretical impossibility of complete verification of ICT systems in the general case is proven for all time. This means that *the best we can ever hope for is best effort quality assurance through testing, with no guarantee that all errors will be exposed.*

Conclusions and Discussion

Computer software, which has been called the most complex human artefact we know, is embedded in complex ICT systems, which in turn are embedded in the larger and even

more complex socio-technical context of human organisations, procedures, operators and users within the built and natural environment. Our ICT systems increasingly pervade the physical and natural environment from aerospace down to nano-scale and beyond; we have instrumented space, the oceans, the earth, our homes, workplaces, vehicles, possessions, our animals and our own bodies.

Our future mHealth systems will be complex distributed ICT systems which are likely to use and interact with emergent ICT and biomedical technologies including smart environments, smart implants, smart drugs, intra-body networks and bionanotechnology.

Furthermore, the complexity of computer systems, communications networks and the amount of data to be transferred and processed is about to increase massively due to several factors. More and more embedded sensor networks generate more and more high volume data (the 'big data' issue). Network complexity of the Internet will increase by orders of magnitude with IPV6. The current IPV4 (the Internet protocol which underlies all Internet communications) has an address space of 2^{32} (meaning that about 4.3 billion devices could be connected to the Internet at any one time). Incredibly, this address space is running out. To cater for future needs, IPV4 is in the process of being replaced by IPV6, with an address space of 2^{128} (allowing about 340 trillion, trillion, trillion devices to be connected simultaneously). This enables an explosive increase in the complexity of Internet topology, in terms of number of nodes and the increase in the potential number of interconnections between nodes. IPV6 thus enables realisation of the Internet of Things (IoT) where any number of everyday objects can have their own Internet address and communicate directly with each other, and can also be directly addressed and even controlled by other nodes. For body area networks, for example, this means that every person and animal in the world could have a sophisticated BAN equipped with innumerable sensors and other devices, each of which could be directly addressable instead of communicating via a platform such as a smartphone.

We can imagine a future *disappearing BAN* enabled (in the short term) by wearable electronics and contactless sensing; in the medium term by implants and in-body wireless communications. Farther into the future we can envisage in-body nano-level sensing, data communications and processing based on cutting-edge research in (bio)nanotechnologies. These developments could enable routine care for large sections of the population where individuals are equipped with nanoBANs, each of which monitors the individual's health condition(s) and delivers *personalised treatment interventions* as well as monitoring *at the cellular level within the body*.

At the same time there will be more and more pressure for future mHealth systems to provide semi-autonomous or autonomous services. So the question of how to quality assure future eHealth and mhealth systems and ensure patient safety in face of the fact that complex ICT systems can exhibit (unintended) emergent behaviour is a crucial one.

The tentative conclusion is that emergencies are emergent phenomena, but probably weakly emergent. ICT systems are both part of the solution and part of the problem. ICT systems themselves, which are now part of the critical infrastructure, carry their own risks and vulnerabilities in emergency response and sometimes exhibit emergent behaviour. We have not succeeded in proving that emergent behaviour in complex ICT systems is a case of strong emergence, however even weakly emergent behaviour poses great challenges. Reductionist explanations, and therefore prediction, seem theoretically possible; however prediction nevertheless seems to be made intractable by the inherent complexity of the systems we have created. That complexity is set to increase explosively for the foreseeable future. In any case, the theoretical impossibility of complete verification of ICT systems in the general case is proven for all time. This means

that the best we can ever hope for is best effort quality assurance through testing, with no guarantee that all errors will be exposed. This fact is well known in computer science circles, but its implications need to be understood by society at large.



Acknowledgements

Grateful thanks are due to Durham University's Institute of Advanced Study for supporting this research under their IAS Fellowship Programme; to my collaborators at Durham University, especially Professor David Budgen, the Institute of Hazard, Risk and Resilience, and to County Durham and Darlington Fire and Rescue Service, Durham Constabulary and Cleveland Fire Brigade for collaboration on ICT for emergency services; to St Cuthbert's Society for their admirable hospitality and welcoming me into collegiate life and to the University of Twente for granting me sabbatical leave to take up the research fellowship at the IAS. The MobiGuide project (<http://www.mobiguide-project.eu/>) has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 287811. The Major Incident Scenario was developed as part of the work of the IST project MOSAIC (FP6-IST-2003-2 004341), funded by the European Commission, and in connection with the Aml@Work initiative of the New Working Environments Unit of the European Commission.

Reference List

Anderson, B. (2012) Emergency futures. *Insights E-Journal* 5(5). University of Durham: Institute of Advanced Study. <https://www.dur.ac.uk/ias/insights/volume5/article5/>

Chalmers, D. J. (200) Strong and Weak Emergence. In Clayton, P. and Davies, P. (eds.) *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion*. Oxford: Oxford University Press.

Department for Communities and Local Government. Fire and rescue national framework for England. First published 11 July 2012. Last updated 15 December 2014. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/5904/nationalframework.pdf

Gorniak, S. (ed.) (2010) Priorities for research on current and emerging network technologies. ENISA (European Union Agency for Network and Information Security) April 20. <https://www.enisa.europa.eu/activities/identity-and-trust/library/deliverables/procent>

HMSO. Civil Contingencies Act 2004. Crown Copyright © 2004. http://www.legislation.gov.uk/ukpga/2004/36/pdfs/ukpga_20040036_en.pdf

JESIP (Joint Emergency Services Interoperability Programme). <http://www.jesip.org.uk/>

Jones, V., Bults, R., Konstantas, D. and Vierhout, P. (2001) Healthcare PANs: personal area networks for trauma care and home care. Fourth International Symposium on Wireless Personal Multimedia Communications (WPMC). Aalborg, Denmark.

Jones, V., Karagiannis, G. and Heemstra de Groot, S. (2005) Ad hoc networking and ambient intelligence to support future disaster response. Proceedings of the IEEE ASWN 2005, 5th Workshop on Applications and Services in Wireless Networks, 137–46.

---. Jones, V., Karagiannis, G. and Heemstra de Groot, S. (2011). Support for Resilient Communications in Future Disaster Management. In Gelenbe, E., Lent, R. and Sakellari, G. (eds.) *Computer and Information Sciences II: 26th International Symposium on Computer and Information Sciences*. London: Springer Verlag, pp. 355–9.

Jones, V. and Saranummi, N. (2010) MOSAIC vision and scenarios for mobile collaborative work related to health and wellbeing. Proceedings of the 1st AMI@work Communities Forum Day 2005, 'Towards ambient intelligence at work 2010', University BW Munich, Germany (22 June, Nottingham University Business School, 108–22. <http://eprints.eemcs.utwente.nl/7458/> (accessed on 30 March 2015).

McConnell, S. (2004) *Code Complete*, 2nd Edition. Redmond, WA: Microsoft Press.

MobiGuide Project: <http://www.mobiguide-project.eu/>

MobiHealth Project: <http://www.mobihealth.org/> 19

Pollock, K. (2013) Review of persistent lessons identified relating to interoperability from emergencies and major incidents since 1986. Emergency Planning College Occasional Papers New Series Number 6. A report commissioned by the Cabinet Office Civil Contingencies Secretariat, October. <http://www.jesip.org.uk/wp-content/uploads/2013/07/Pollock-Review-Oct-2013.pdf>

Report of the 7 July Review Committee Volume 4: Follow-up Report, London Assembly, August 2007. legacy.london.gov.uk/assembly/reports/7july/follow-up-report.rtf (accessed on 30 March 2015).

The 9/11 Commission Report. Final report of the National Commission on Terrorist Attacks upon the United States. Executive summary. <http://www.9-11commission.gov/report/index.htm> (accessed on 30 March 2015).

Turing, A. (1937) On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society, Series 2, Volume 42, pp. 230–65. DOI:10.1112/plms/s2-42.1.230.

Wood-Heath, M. and Annis, M. (2004) Working together to support individuals in an emergency or disaster. British Red Cross, May. <http://www.cabinetoffice.gov.uk/resource-library/working-together-support-individuals-emergency-or-disaster> (accessed on 30 March 2015).

Zoraster, R. M. (2006) Barriers to disaster coordination: health sector coordination in Banda Aceh following the South Asia Tsunami. *Prehospital and Disaster Medicine* 21(1): 13–18, February. Copyright © World Association for Disaster and Emergency Medicine 2006. DOI:[http:// dx.doi.org/10.1017/S1049023X0001582X](http://dx.doi.org/10.1017/S1049023X0001582X)

Backlist of Papers Published in Insights

No.	Author	Title	Series
2008 Volume 1			
1	Boris Wiseman	Lévi-Strauss, Caduveo Body Painting and the Readymade: Thinking Borderlines	General
2	John Hedley Brooke	Can Scientific Discovery be a Religious Experience?	Darwin's Legacy
3	Bryan R. Cullen	Rapid and Ongoing Darwinian Selection of the Human Genome	Darwin's Legacy
4	Penelope Deutscher	Women, Animality, Immunity – and the Slave of the Slave	Darwin's Legacy
5	Martin Harwit	The Growth of Astrophysical Understanding	Modelling
6	Donald MacKenzie	Making Things the Same: Gases, Emission Rights and the Politics of Carbon Markets	Modelling
7	Lorraine Code	Thinking Ecologically about Biology	Darwin's Legacy
8	Eric Winsberg	A Function for Fictions: Expanding the Scope of Science	Modelling
9	Willard Bohn	Visual Poetry in France after Apollinaire	Modelling
10	Robert A. Skipper Jr	R. A. Fisher and the Origins of Random Drift	Darwin's Legacy
11	Nancy Cartwright	Models: Parables v Fables	Modelling
12	Atholl Anderson	Problems of the 'Traditionalist' Model of Long-Distance Polynesian Voyaging	Modelling
2009 Volume 2			
1	Robert A. Walker	Where Species Begin: Structure, Organization and Stability in Biological Membranes and Model Membrane Systems	Darwin's Legacy
2	Michael Pryke	'What is Going On?' Seeking Visual Cues Amongst the Flows of Global Finance	Modelling
3	Ronaldo I. Borja	Landslides and Debris Flow Induced by Rainfall	Modelling
4	Roland Fletcher	Low-Density, Agrarian-Based Urbanism: A Comparative View	Modelling
5	Paul Ormerod	21st Century Economics	Modelling
6	Peter C. Matthews	Guiding the Engineering Process: Path of Least Resistance versus Creative Fiction	Modelling
7	Bernd Goebel	Anselm's Theory of Universals Reconsidered	Modelling
8	Roger Smith	Locating History in the Human Sciences	Being Human
9	Sonia Kruks	Why Do We Humans Seek Revenge and Should We?	Being Human
10	Mark Turner	Thinking With Feeling	Being Human
11	Christa Davis Acampora	Agonistic Politics and the War on Terror	Being Human
12	Arun Saldanha	So What <i>Is</i> Race?	Being Human
13	Daniel Beunza and David Stark	Devices For Doubt: Models and Reflexivity in Merger Arbitrage	Modelling
14	Robert Hariman	Democratic Stupidity	Being Human

No.	Author	Title	Series
2010 Volume 3			
1	John Haslett and Peter Challenor	Palaeoclimate Histories	Modelling
2	Zoltán Kövecses	Metaphorical Creativity in Discourse	Modelling
3	Maxine Sheets-Johnstone	Strangers, Trust, and Religion: On the Vulnerability of Being Alive	Darwin's Legacy
4	Jill Gordon	On Being Human in Medicine	Being Human
5	Eduardo Mendieta	Political Bestiary: On the Uses of Violence	Being Human
6	Charles Fernyhough	What is it Like to Be a Small Child?	Being Human
7	Maren Stange	Photography and the End of Segregation	Being Human
8	Andy Baker	Water Colour: Processes Affecting Riverine Organic Carbon Concentration	Water
9	Iain Chambers	Maritime Criticism and Lessons from the Sea	Water
10	Christer Bruun	Imperial Power, Legislation, and Water Management in the Roman Empire	Water
11	Chris Brooks	Being Human, Human Rights and Modernity	Being Human
12	Ingo Gildenhard and Andrew Zissos	Metamorphosis - Angles of Approach	Being Human
13	Ezio Todini	A Model for Developing Integrated and Sustainable Energy and Water Resources Strategies	Water
14	Veronica Strang	Water, Culture and Power: Anthropological Perspectives from 'Down Under'	Water
15	Richard Arculus	Water and Volcanism	Water
16	Marilyn Strathern	A Tale of Two Letters: Reflections on Knowledge Conversions	Water
17	Paul Langley	Cause, Condition, Cure: Liquidity in the Global Financial Crisis, 2007–8	Water
18	Stefan Helmreich	Waves	Water
19	Jennifer Terry	The Work of Cultural Memory: Imagining Atlantic Passages in the Literature of the Black Diaspora	Water
20	Monica M. Grady	Does Life on Earth Imply Life on Mars?	Water
21	Ian Wright	Water Worlds	Water
22	Shlomi Dinar, Olivia Odom, Amy McNally, Brian Blankespoor and Pradeep Kurukulasuriya	Climate Change and State Grievances: The Water Resiliency of International River Treaties to Increased Water Variability	Water
23	Robin Findlay Hendry	Science and Everyday Life: Water vs H ₂ O	Water
2011 Volume 4			
1	Stewart Clegg	The Futures of Bureaucracy?	Futures
2	Henrietta Mondry	Genetic Wars: The Future in Eurasianist Fiction of Aleksandr Prokhanov	Futures
3	Barbara Graziosi	The Iliad: Configurations of the Future	Futures
4	Jonathon Porritt	Scarcity and Sustainability in Utopia	Futures
5	Andrew Crumey	Can Novelists Predict the Future?	Futures
6	Russell Jacoby	The Future of Utopia	Futures
7	Frances Bartkowski	All That is Plastic... Patricia Piccinini's Kinship Network	Being Human

No.	Author	Title	Series
8	Mary Carruthers	The Mosque That Wasn't: A Study in Social Memory Making	Futures
9	Andrew Pickering	Ontological Politics: Realism and Agency in Science, Technology and Art	Futures
10	Kathryn Banks	Prophecy and Literature	Futures
11	Barbara Adam	Towards a Twenty-First-Century Sociological Engagement with the Future	Futures
12	Andrew Crumey and Mikhail Epstein	A Dialogue on Creative Thinking and the Future of the Humanities	Futures
13	Mikhail Epstein	On the Future of the Humanities	Futures

2012 Volume 5

1	Elizabeth Archibald	Bathing, Beauty and Christianity in the Middle Ages	Futures II
2	Fabio Zampieri	The Holistic Approach of Evolutionary Medicine: An Epistemological Analysis	Futures II
3	Lynnette Leidy Sievert	Choosing the Gold Standard: Subjective Report vs Physiological Measure	Futures II
4	Elizabeth Edwards	Photography, Survey and the Desire for 'History'	Futures II
5	Ben Anderson	Emergency Futures	Futures
6	Pier Paolo Saviotti	Are There Discontinuities in Economic Development?	Futures II
7	Sander L. Gilman	'Stand Up Straight': Notes Toward a History of Posture	Futures II
8	Meredith Lloyd-Evans	Limitations and Liberations	Futures II

2013 Volume 6

1	David Martin-Jones	The Cinematic Temporalities of Modernity: Deleuze, Quijano and <i>How Tasty was my Little Frenchman</i>	Time
2	Robert Levine	Time Use, Happiness and Implications for Social Policy: A Report to the United Nations	Time
3	Andy Wood	Popular Senses of Time and Place in Tudor and Stuart England	Time
4	Robert Hannah	From Here to the Hereafter: 'Genesis' and 'Apogenesis' in Ancient Philosophy and Architecture	Time
5	Alia Al-Saji	Too Late: Racialized Time and the Closure of the Past	Time
6	Simon Prosser	Is there a 'Specious Present'?	Time

2014 Volume 7

1	Robert Fosbury	Colours from Earth	Light
2	Mary Manjikian	Thinking about Crisis, Thinking about Emergency	Time
3	Tim Edensor	The Potentialities of Light Festivals	Light
4	Angharad Closs Stephens	National and Urban Ways of Seeing	Light
5	Robert de Mello Koch	From Field Theory to Spacetime Using Permutations	Time
6	Jonathan Ben-Dov	What's In A Year? An Incomplete Study on the Notion of Completeness	Time

No.	Author	Title	Series
7	Lesley Chamberlain	Clarifying the Enlightenment	Light
8	Fokko Jan Dijksterhuis	Matters of Light. Ways of Knowing in Enlightened Optics	Light
9	Paul O'Brien	Understanding Nano-Science and -Technology: Towards a Definition of a Nanocrystal	Light

Insights

Insights is edited by Barbara Graziosi, IAS Director and Professor of Classics.
Correspondence should be directed to Pauline Edmondson (pauline.edmondson@durham.ac.uk)