

Simultaneous trend analysis for evaluating outcomes in patient-centred health monitoring services

Edward C. Conley · David R. Owens ·
Stephen L. Luzio · Mahesh Subramanian ·
Ali Shaikh Ali · Alex Hardisty · Omer Rana

Received: 18 July 2007 / Accepted: 14 February 2008 / Published online: 22 April 2008
© Springer Science + Business Media, LLC 2008

Abstract The research aim underpinning the *Healthcare@Home (HH)* information system described here was to enable ‘near real time’ risk analysis for disease early detection and prevention. To this end, we are implementing a family of prototype web services to ‘push’ or ‘pull’ individual’s health-related data *via* an system of clinical hubs, mobile communication devices and/or dedicated home-based network computers. We are examining more efficient methods for ethical use of such data in timeline-based (i.e. ‘longitudinal’) data analysis systems. A consistent data collation infrastructure is being created for use along the ‘patient path’—accessible wherever patients happen to be. This ‘patient-centred’ infrastructure can be applied in the evaluation of disease progression risk (in the light of clinical understanding of disease processes). In this paper we describe the requirements for making multi-data trend management ‘scale-up’, together with some requirements of an ‘end-to-end’ functioning data collection system. A Service-Oriented Architecture (SOA) approach is used to maximise benefits from (1) clinical evidence and (2) computational models of disease progression that can be made available elsewhere on the SOA. We discuss the

implications of this so-called ‘closed loop’ approach for improving healthcare intervention outcomes, patient safety, decision support, objective measurement of service quality and in providing inputs for quantitative healthcare (predictive) modelling.

Keywords Web services · Time series analysis routines · Scalability · Chronic disease management · Portal technologies · Risk monitoring · Service-oriented architecture

1 Introduction—building a scaleable information system for early detection and prevention of disease

We are interested in developing *interoperable* families of web-based information services for ‘near-real time’ health risk factor data collation and multiple trend and *time series analysis*. These information services are required to more consistently evaluate and model disease progression risk—both at the individual level and for *aggregate* population-based reporting. The approach needs a consistent longitudinal (timeline-based) *analytical framework* [1] that is also applicable to ethical research e.g. for comparative analysis of intervention impacts and prediction of outcomes using quantitative modelling [27]. For disease prevention applications, the underlying information services need to be useable at any stage from health, through incipient disease, to severe complications.

In the case of diabetes, this means support for (1) prevention in healthy individuals (2) prevention in the ‘pre-diabetes’ stage (3) support for managing comprehensive long-term risk factor collation and (4) support for monitoring signs and progression of complications such as retinopathy, nephropathy or neuropathy. A significant

To aid understanding, a concise glossary is provided for italicized technical or common ‘jargon’ terms that are not defined in the text.

E. C. Conley (✉) · D. R. Owens · S. L. Luzio
Diabetes Research Unit, School of Medicine, Cardiff University,
1st Floor, Academic Centre, University Hospital Llandough,
Penlan Road, Penarth,
Cardiff CF64 2XX Wales, UK
e-mail: e.c.conley@cs.cardiff.ac.uk

E. C. Conley · M. Subramanian · A. S. Ali · A. Hardisty · O. Rana
The Welsh e-Science Centre,
Cardiff University School of Computer Science,
5 The Parade, Roath,
Cardiff CF24 3AA Wales, UK

barrier for ‘scale-up’ these approaches (for application in disease early detection and prevention across large populations) is inconsistency in data collection across sites. Another confounding factor is the prevalence of non-standardised, ‘home-grown’ or *ad hoc* methods for information exchange. This paper therefore reports work towards a more consistent approach for ‘end-to-end’ risk data collection and computational analysis called *Health care@Home (HH)*. This is an integrated research model for *inter-enterprise computing* applied to patient-centred, disease early detection/prevention where healthcare is patient-centred and contiguous across clinic, home and mobile locations.

2 Improving consistency of data collection along the ‘patient path’ for improved decision support

Amongst its objectives, *HH* creates a systematic integration framework for co-development of advanced healthcare decision support tools through a continuously-operating ‘closed loop’ *outcomes analysis* model. With specific informed consent and ethical approval, system-standards compliant hardware can sample data from anywhere on the ‘patient path’—when in the clinic, when mobile, and when ‘at-home’. *HH* is predicated on (largely unrealised) healthcare benefits promised by application of scalable distributed computing using *grid* and *Service Oriented Architecture (SOA)* technologies (see later and refs. [2, 6, 16, 19, 22]). *HH* has, however avoided the trap of being ‘technology-driven’.

The *HH* design includes an extensible family of composite services through “portlets” (see Section 6) to directly support predefined, ‘domain-driven’ workflows within integrated care pathways (ICP’s). The portlets (and underlying information services) are each deliberately limited in scope and complexity. According to a *logical domain model* that satisfies end-user need, all services integrate via *portal* technology to deliver ‘*composite process functions*’. This information model-driven integration (with potential for re-use of underlying standardised services and devices) can make risk data gathering support more consistent along constantly changing patient paths – which are likely to include mobile and home applications.

3 A fundamental role for ‘domain-driven’ business process modeling

In the *Healthcare@Home* project, *integrated care pathways* (ICP’s) provided some of the underlying ‘domain’ knowledge and business process logic for information service co-integration. ICP’s used in the project were associated with

the Wales National Service Framework (NSF) for Diabetes [26] which provided a critical set of non-redundant workflow processes—in this case ‘real-world’ clinical practices requiring data support. The stakeholder integration value of the NSF was recognised early on. It certainly represented an authoritative policy statement that healthcare professionals would look to for setting out forward-looking consensus practice standards (assisting patient-centred, proactive management of conditions).

The logic in the NSF therefore became the foundation for the supporting information system structure. This ‘domain-driven’ approach, characterised the project to some as a ‘bottom-up’ *policy deployment mechanism* for achieving the holistic (disease prevention) aims of the NSF. The shared purpose principle also helped gain clinical practitioner and patient acceptance, understanding and ‘ownership’ of the demonstrator solutions produced. This approach was in contrast to so-called ‘top-down’ approaches that apply a standard solution without detailed understanding of behaviours and needs within the domain itself. Involvement of real end-users and *domain knowledge experts* (i.e. who have deep understanding of the domain and its complexities) has been shown to be critical for building a ‘pipeline’ of interacting information services.

Special efforts were made to understand the domain and capture this understanding in consistent models of information flow. Despite this, the more systematic approach known as *domain-driven design* [32] has not been applied in the project to date, but remains an interesting possibility for future developments. A more structured, process-based approach to information service development will be necessary where larger numbers of practitioner-stakeholders are involved and where software solutions need to ‘scale’. More systematic methods for capturing ‘deep understanding’ of real-world business practices will therefore supplement simple ‘bottom-up’ analysis of needs. Understanding of how disease prevention can be practiced more effectively (driven by information flow in many different circumstances along the patient path) will need to be captured in consistent logical domain models. Especially complex domain issues e.g. the rules of clinical information governance, the ethics of data use/re-use and patient privacy concerns might benefit from formal business process modelling activities, and will be essential for building ‘fit-to-need’ information services.

4 Building-in ethical safeguards pertaining to data use and re-use

The currently-developed research-phase of the *HH e2e* data gathering is also focused on improved multi-trends management. A future goal is to use the framework for more

robust approaches to *algorithmic modelling of outcomes* using data collected in ‘near real time’. Such computational analyses and ‘dynamic graphing’ of individual’s data trends are intended to be the groundwork for data monitoring services in a way that is comprehensible to end-users. These technical solutions are constrained by the requirements for rigorous and unambiguous ethical policies on data reuse—in other words a well-researched and applied *information governance* policy.

As sensitive data is more readily exchanged and aggregated through adoption of technical interoperability standards, we believe more stringent ethical and privacy safeguards need to be in place to prevent data misuse. ‘Misuse’ might include uses of data that were not part of the informed consent process. The same considerations need to be given in data reuse by quantitative modelling approaches using (for example) longitudinal statistical analysis, algorithmic programming, outcome risk modelling and system simulations. All potential enduser communities (including those specifically interested in quantitative modelling using service-derived data, [27]) need to act collectively to guarantee adoption of ethical safeguards when data is aggregated in information services with predictive functions.

5 *Healthcare@Home* as a ‘closed-loop’ disease prevention/management system

Generic disease early detection and prevention services are a development priority the world over. The need for consistent data and information management services—to address the overwhelming global burden of disease—remains the *raison d’être*. We see the development of consistency through data interoperability as having both local-action and global-action components. While a quest for *global* interoperability remains largely unrealised, it can be a key driver for achievement of economies of scale, sustainability of large-scale medical device/information service systems. Wider (if not universal) deployment of consistent information systems/device families will be necessary for maximising unified approaches to disease prevention—irrespective of deployment site.

HH is characterised as an ‘end-to-end/anywhere-to-anywhere’ (e2e/a2a) information integration system that functions to bring different aspects of open standards-based technology together to work for a *common purpose* of disease early detection/prevention. In working out how this *HH* ‘e2e’ might function, real-world ‘clinical, at home and mobile’ demands for information were analysed to design a core set of composite information services. In their simplified form (Fig. 1, upper) the *HH* ‘core closed loop’ comprises composite information services supporting effi-

cient data acquisition (Fig. 1, step 1, implementable as a streamlined or ‘lean’ clinical workflow) that can simultaneously send (bind together) patient-identifier and device-type outcome risk data from a prescribed set of instruments (or patient history forms) into a clinical data hub integrator device *en route* to the decision support portlets and the patient’s own electronic health records (EHR).

In the *HH* workflows, high quality clinical assessment data is routinely assessed for ‘individualised risk’ (Fig. 1, step 2) to prioritise patients in need of urgent referral and to compare their status with a ‘risk signature’ that might indicate that an immediate preventative or screening action needs to be taken. As a result of this initial analysis, ordering of a number of optional specialised diagnostic tests (Fig. 1, step 3) may be triggered. Individualised risk variables are then visualised by forward-population of a ‘relative risk dashboard’ tool (Fig. 1, step 4) to help the clinician judge the status of the condition’s risk indicators against ranges observed in the population. The relative risk dashboard therefore provides visual context and support for taking the best stepwise options for the individual’s care and treatment.

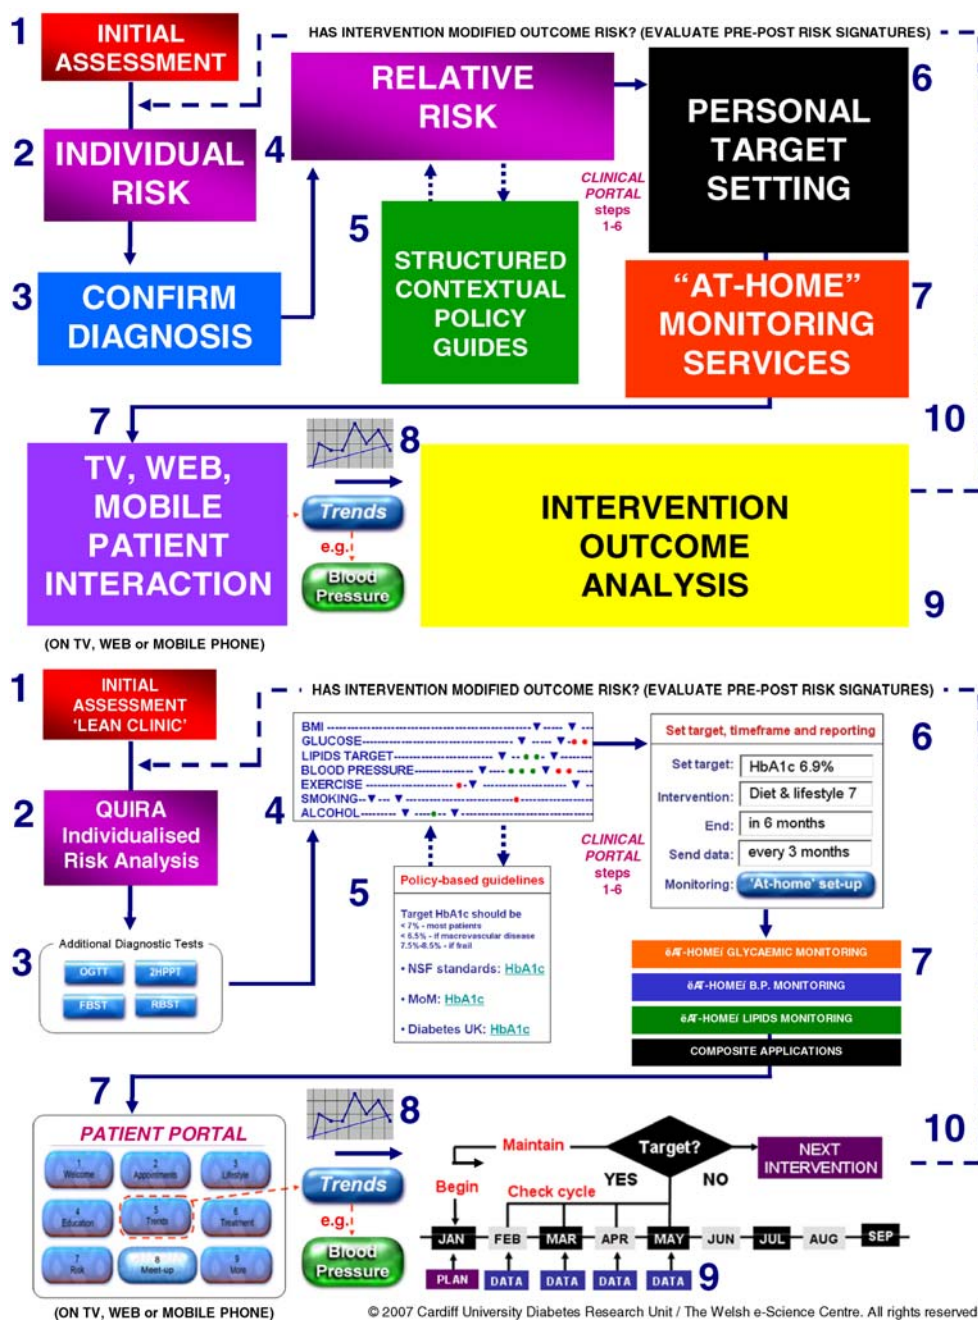
6 Mapping to individualised care planning, stratified risk and intervention analysis services

More detailed policy-based guidelines affecting the domain can be ‘pulled’ or ‘pushed’ at the most relevant points in the decision support workflow (i.e. from other SOA information services by relevance-of-fit to clinical options at that ‘decision point’—Fig. 1, step 5). Following face-to-face discussion of those options with the patient (a critical step aimed at giving the patient active involvement in his/her episode of care), a specific intervention may be selected from options that are supported by evidence of efficacy.

If appropriate, a personal outcome ‘target’ can be set along with other compliance requirements—Fig. 1, step 6. i.e. Personal target setting as part of the pathway requires active engagement/participation of the patient - in line with the principles of the NSF. With personal participation and informed consent taken on the use of data that will be generated *within* the specific monitoring episode, a selection of composite ‘at home’ applications that themselves work with the patient’s monitoring devices (e.g. *via* IPTV, PC or phone portal clients, Fig. 1, step 7 and Fig. 6) are invoked.

Coded data flowing between the patient’s and clinical team’s display portals should efficiently manage an entire intervention ‘episode’ and be capable of processing incoming data from specialised but standardised instrument types. *Prescriptive* choices by the healthcare professional (choices made as a matter of judgement) write ‘procedural’

Fig. 1 *Upper:* Simplified representation of the ‘end-to-end, closed-loop’ system of composite information services associated with *Healthcare@Home*. The figure reflects both the logical workflow recognised by clinical process and the ‘division of labour’ between different computational services that interact around a web services integration platform or messaging fabric in an SOA. The framework is extensible to support entire integrated care pathways underpinned by NSF policies. The steps in the figure are outlined in the text. *Lower:* More detailed representation of the workflow indicating some key features. Aspects described in the text include QUIRA (*QU*antitative *I*ndividualised *R*isk *A*nalysis, *s*tep 2); the relative risk dashboard (*s*tep 4), the ‘at home’ monitoring services combined with the IPTV/PC/mobile patient portal (*s*tep 7), timeline-based multi-trends services (*s*tep 8) and the ‘intervention efficacy’ or outcome analysis engine (*s*tep 9)



content into detailed individualised care plans, and automatically generate appointment messages that can be consumed by a scheduling service. At this point, different members of the care team (through their role-based view of the clinical portal) may be alerted, bringing additional professional support pertinent to the type of care package being delivered or other treatment/intervention. Progress is monitored systematically by an intervention outcomes analysis application (Fig. 1, step 8, based on a longitudinal, multi-trend database design using a single dataset between all participating sites).

7 Providing support for integrated care pathways

This converges multi-modality data on to one timeline to track not only multiple (physiological) trends but also other clinical/life events that are pertinent. The ‘evaluation of outcomes’ is another critical part of the closed loop management cycle, and is further elaborated in Fig. 2. Additional workflow (not shown) captures any deviations from the plan towards the agreed personal target (Fig. 1, step 9). Automating the recording of steps in the ICP in this way is intended to capture business process criteria

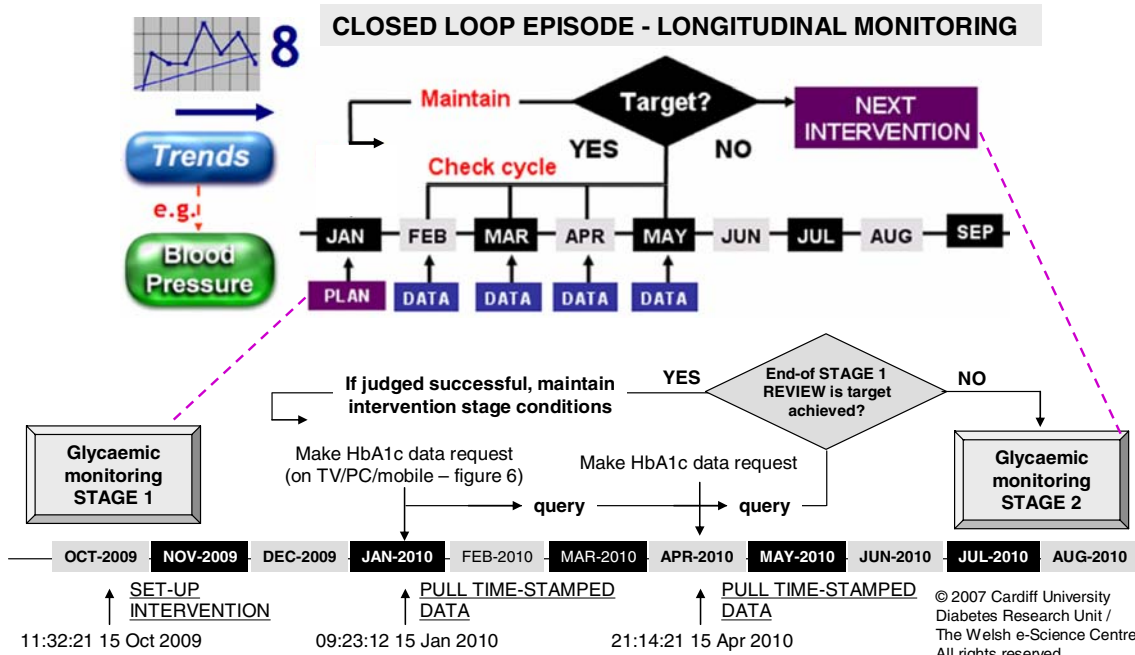


Fig. 2 Figurative connections between ‘at-home and mobile’ data collection, trends services and iterative ‘algorithmic’ intervention monitoring. The personal target setting cycles of the Diabetes NSF workflows specify the algorithm for longitudinal data gathering. This includes active querying of ‘intervention outcome’ as part of case management e.g. for HbA1c data. Through an SOA approach this would invoke and coordinate patient identity, blood sample tracking and laboratory results services, returning quantitative results to clinical and patient portlets. Note the figures do not reflect underlying functionality, not the user interface. Features include intuitive visual

timelining (in the portals, see Fig. 6), multi-event-type symbolic marking (*ibid.*) and pathway stage exit/entry following ‘end-of-stage’ reviews. The key management benefit is that staging and ‘near real time status’ is immediately apparent. The ‘closed loop’ relates to the ‘algorithmic’ nature of the NSF-driven pathways—decisions are taken on the basis of *evaluated data* arriving via the monitoring services *on a continuum* starting with the comprehensive initial assessment generating quality-controlled data (see Fig. 1, step 1). Figure 5 explores how this iterative closed loop risk monitoring principle can be made ‘scaleable’

determined by an outcome analytical algorithm—e.g. whether (given conditions XXZ) patients will revisit the clinic or ‘progress’ to the next part of the care pathway.

The ‘closed loop’ feedback aspect of the workflow is represented by the dotted line (Fig. 1, step 10) to the individualised risk service called ‘QUIRA’. The closed loop acts as a ‘pre- to post-intervention’ reference point for *relative intervention validity and efficacy*—i.e. whether a prescribed action was ‘performed to plan’ and if it was, whether it ‘worked or not’ as judged by independent performance criteria. This continuous workflow system also has a key advantage that patient status can be checked in *near-real time* whilst being flexible enough to challenge barriers to self-management.

8 Optimising disease prevention through a ‘closed loop’ system

Improved data quality enabled by data definition and device interoperability standardisation will provide a more robust basis for clinical judgement. With reference to the *HH* workflow, each block in the core ‘closed analysis loop’

(described by Figs. 1 and 2) closely reflects a part of the domain (Diabetes NSF) policy and can thus help with its deployment at larger scale. The scope of *HH* permits incorporation (and exchange) of standardised multiple modality device types, standardised ‘patient path’ data transducer/router devices (e.g. data ‘hubs’ that can securely communicate data to a person’s xHR) and standardised (semantically constrained) web services. Program components can be functionally diverse and brought together (‘called’ or ‘discovered’) as part of an open-standards SOA using a harmonised messaging fabric or commercial integrator platform (*HH* has initially been implemented in IBM’s *Websphere* product suite, see Fig. 5).

An SOA approach can enable specialised data processing services to access common clinical information services subject to the ethical (data re-use) constraints and permissions enabled by the patient at step 1 of the workflow. In Fig. 1, large granularity services have been aligned to pragmatic tasks that make sense to health care professionals (Fig. 1 upper). Ideally they contribute business process value by improving routine (repeated) data gathering while improving workflow efficiency/safety/data quality. Services can also generate ‘process value’ by using a *clinical*

outcome-based integration principle to constrain implementation choices. In general terms, this is needed because of the overwhelming complexity of the healthcare technology domain (combined with massive ‘optionality’ offered by global data device ecosystems). ‘Too much optionality’ in technology implementation choices generally creates increased complexity - and as a consequence can introduce unnecessary patient risk or development delays [3]. Teams of healthcare professionals need efficient information systems to help them comply with a complex collection of local/national role-specific policies. *This complexity control can help them gain relevant knowledge to apply locally-available ‘best practice’ solutions.*

9 Choice of interoperability standards

Historically (and notably in the field of diabetes) multiple incompatibilities between non-integrated, proprietary blood glucose sensors, device types and software applications have forced information service providers to adopt multiple non-standardised ‘point-to-point’ models of patient data monitoring. This non-standardisation has led to massive cost increases, wastage of service resources, duplication of incompatible device types, inefficient spending of training resources, early obsolescence of equipment and significantly, poor quality knowledge representation (with non-optimal ethical reuse of data). A purpose-led, policy-mapped prioritised model like *HH* can constrain unnecessary complexity generated within an ‘end-to-end’ ecosystem introduced by vendor component optionalities (which in isolation may be innocuous, attractive, market-differentiating features). The system view can also prevent semantic inconsistencies at the user interface level which might lead to patient safety concerns.

Without a single system view, problems can arise by ‘emergent’ properties when different untested combinations of system components come together in multiple global settings. For its development phase, *HH* is now considering adoption of a number of harmonisation frameworks, standards and other innovations that can improve end-to-end system consistency. To be effective they also must help the ‘anywhere-to-anywhere’ (‘a2a’) goal by being practiced/promoted across many countries and thus increasing in their ‘momentum’ for enduser adoption. Ideally, these frameworks will be converging to a ‘global’ integration model (examples might include CDA Release 2, IHE/XDS/XDS-I [7, 31], SNOMED-CT [9], and for point-of-care applications, the Continua Healthcare Alliance (CHA, [24]) and CEN/ISO/IEEE (→ IEEE11073). Other national-level innovations like the Common User Interface (CUI, [29]) and established consensus domain datasets like the DCCR (see next section) may also play an important role, as will

any coherent, complementary datasets (e.g. for diabetic retinopathy data).

10 Dataset harmonisation across multiple sites

Achievement of a ‘disease early detection and prevention system’ meeting patient-centred operational requirements demanded by the Diabetes NSF requires technical interoperability to be an essential first step. The preservation of *meaning* across complex health information systems is a far greater challenge. This is not only because clinical words/concepts can easily be confused in search terms (multi-meanings of the same term, similar meanings of different terms, multi-terms for same meaning) but also change meaning through ‘cumulative error’ as non-standardised information flows through networks.

A better system would avoid data transformations and use information that is intelligible to different core and specialist services—in *addition to being human readable*. Of the innovations described above, CDA Release 2 [8] is now delivering capability in this function. Much research is also dedicated to the so-called ‘semantic web’ that may underpin and enable next-generation services that re-use data (subject to ethical constraints mentioned earlier). Such approaches can be enhanced if a consistent e2e system policy exists for meaning (semantics) at points of data collection, entry, consumption and use. This is easier said than done, especially with the significant conceptual complexity inherent to the patient-centred, multi-disciplinary healthcare process. Semantic interoperability issues also profoundly affect the power of quantitative modelling applications as highlighted in a later section.

There are many aspects to semantic interoperability (most beyond the scope of this paper) but in basic terms, one of the most essential is trans-system dataset harmonization—the agreed set of data definitions, units of measurement, precise definitions and messaging standards that permit each site to ‘push’ and ‘pull’ information—avoiding the ‘cumulative error’ problem. Datasets thus comprise structured lists of individual data items, each with a clear label, definition and set of permissible values, codes and classifications. From this, subject to ethical and data protection constraints, secondary uses of information can be derived or compiled or both in support of specific health and social care purposes. However, comprehensive datasets and terminologies are notoriously difficult to create—they are typically evolved from overlapping legacy datasets or terminologies on a ‘consensus by committee’ involving a non-streamlined voting/approvals/maintenance process. However, once in existence, domain datasets/terminologies can be an important contribution (but not a solution) to semantic interoperability. For the *HH* project, adoption of a

single domain-driven dataset was a cornerstone for producing a single source of risk data definitions produced by different system inputs (devices, hubs, electronic form-based history pick-lists etc.). These definitions provided a means to metadata tag and aggregate data outputs (in near real time or historically from different sites).

11 The diabetes continuing care reference dataset

For starting to build its individualised risk analysis framework, the *HH* project adopted the Diabetes Continuing Care Reference dataset (DCCR, see [25]) which is designed to support the delivery of the Diabetes NSF across a range of different care settings. The DCCR specification combines data monitoring requirements of four diabetes work streams previously practiced in the UK: the National Diabetes Audit (NCASP); the diabetes chapter of nGMS QoF, the Diabetes e-Performance Management Tool and the diabetes indicators for the Better Metrics Performance Indicator Project. Datasets like DCCR however are key for building consistent enduser combinatorial queries (i.e. ‘declarative’ queries that are *scriptable*). The choice of the DCCR still does not solve inconsistencies in data definitions on a ‘global’ basis (e.g. non-standardised units, test reporting formats) and so these remain sources of potential error in development of an ‘a2a’ interoperable risk framework. The relationship of the DCCR to the *HH* ‘at home and mobile’ trends monitoring services and other datasets supporting the NSF’s are shown in Fig. 3. An investment in datasets is fundamental to improving content labelling, accessibility, structured search that underpin planning, delivery and monitoring of services within health and social care. As described by NHS Information Centre, datasets are also intended to support information requirements of national and local performance management, planning, commissioning and clinical governance, quality assurance in the monitoring of National Service Frameworks (NSFs). Fundamentally, they can assist same standard of information to be generated from care records, independent of the organisation or system that captures the base data.

12 Interfacing scaleable SOA monitoring services, ‘HealthGrids’ and ‘data device ecosystems’

The core *HH* project will implement a family of prototype web services that support multiple trend services that are relevant along the ‘patient path’ (i.e. to ‘push’ or ‘pull’ individual’s health-related data from dedicated clinical and/or home-based network servers or mobile communications devices to one or more data analysis engines). Our principal

interests are in developing ‘near-real time’ data collation and time series analysis routines required to evaluate the status of individual’s condition and in using a disease progression model to understand and compute ‘progression risk’ for the development of complications. One key challenge is in optimising the approach so it can be ‘scaled’ safely and robustly to monitor large numbers of patients with multiple conditions e.g. in chronic disease management services. Server-side algorithmic analyses and dynamic graphing of individual trend data outputs are available to the patient (or ethical monitoring services) through a timeline-based portal exploiting a flexible approach that is comprehensible to end-users. The model is designed to monitor a wide range of online risk data from sensor devices within home or mobile environments.

As indicated in the datasets section, it was important to support legitimate secondary uses through data aggregation functions—either on ‘dashboards’ that could be made available for public health research purposes or for quantitative healthcare modelling applications. To date we have only researched the operational requirements for improved trends management and defined the logical and physical architecture requirements of the system (next section). A Service-Oriented Architecture (SOA) approach is being used to maximise interoperability with articulated clinical evidence or with computational models of disease progression (either exposed as consumable information services elsewhere on the SOA). Such approaches form part of ‘next generation’ promise for ubiquitous resource discovery in ‘HealthGrids’ that make available a wide range of web service technologies in a scaleable manner [28] and refs. [4, 5, 10, 13–15, 20, 23]. HealthGrids rely on standards-based services (e.g. for storage or retrieval of data in repositories) and are likely to have an increasing role supporting ethical secondary uses including epidemiological research applications and the normal range of clinical functions in the *HH* e2e (Fig. 1).

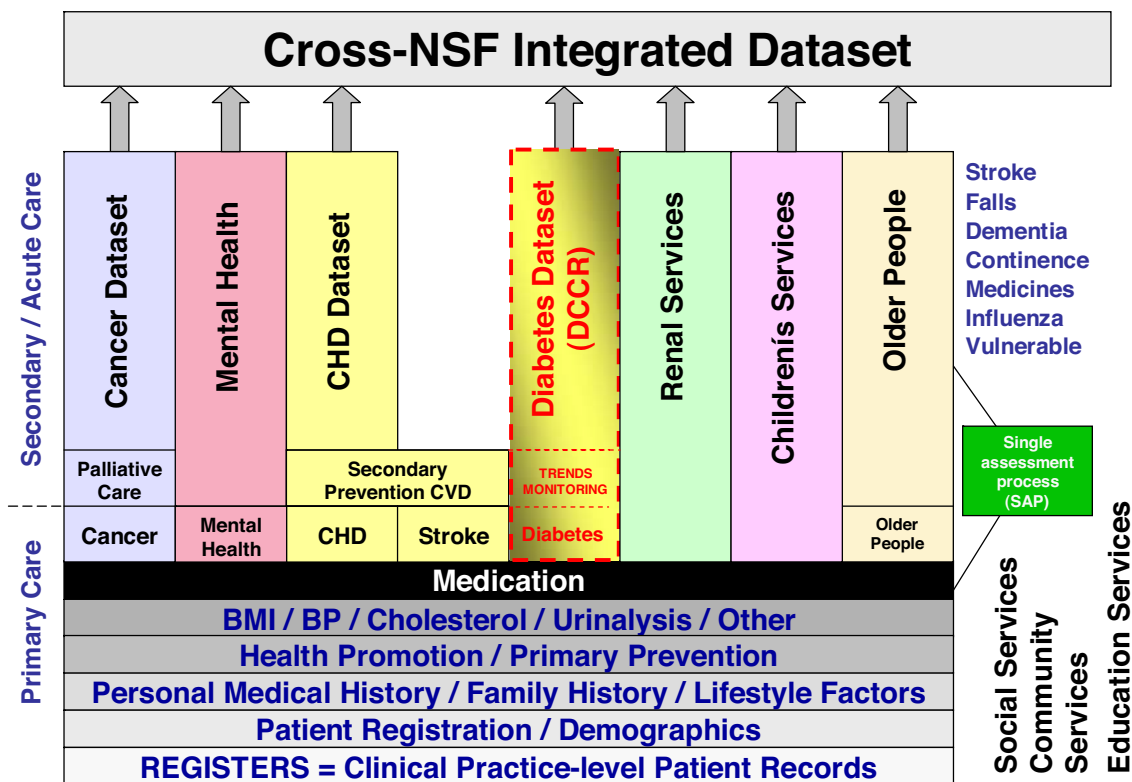
13 Use of the *HH* framework in clinical outcomes research

To demonstrate its applications in clinical outcomes research, *HH* has implemented a limited range of Boolean (combinatorial) search functions that nevertheless can retrieve named data items according to the ‘ubiquitous language’ of the DCCR template. As exemplified by the origins of *HH* in the Diabetes National Service Framework, we would re-assert that HealthGrids will ‘fit-to-need’ best if they are built ‘bottom-up’ closely following the behaviours of the domain. Many notable HealthGrid projects in a variety of specialised application areas are reviewed and consolidated by annual HealthGrid meetings. These ‘mod-

ular’ interoperable services (compare Fig. 1 upper) can be predicted to fuel growth of the HealthGrid approach. Broad granularity services will be functionally quite diverse, yet (as exemplified by the *HH e2e*) retain core interoperability. Information services may enable for example, consistent support for ‘patient path’ operations such as personalised monitoring, radiation dose computation, multi trends analysis, standardised data transformations (e.g. statistical treatments), intervention management, treatment regimen, imaging service data acquisition (e.g. MRI, CT, PET, ultrasound, retinal images), database lookups (e.g. pathway-contextualised best practice), database calls (e.g. genetics tests), procedure planning, annotation/terminology/datasets, specialised diagnosis, team conference case work-

flows, data visualisation and quantitative modelling and simulation services.

Common clinical or generic management services may co-exist with these specialised services in HealthGrids, examples being those supporting patient registries, service descriptions (retaining semantic consistency of description for exposure or discovery through use of web services description language or WSDL), configuration/maintenance services, scheduling and longitudinal data services (as exemplified by Fig. 2) and data interface management (e.g. that may assist transfer of data from one repository to another subject to express patient consent on the data re-use). As with all other aspects of the ‘e2e’—reliable discovery of functional resources on the SOA largely



Modified by DRU from figure annotated Michelle Rawle (NHS HSCIC) in 2005. Original Copyright © 2007, re-used with the permission of The Information Centre on behalf of NHS. All rights reserved.

Fig. 3 Strategic dataset adoption in the *Healthcare@Home* ‘end-to-end, anywhere-to-anywhere’ (‘e2e-a2a’) research model for data gathering. Other diabetes datasets in development may become relevant for *HH*. These include the Diabetic Retinopathy Screening Dataset (currently at the approvals process for ISB Full Operational Standard approval) and the National Diabetes Audit Dataset (currently going through the approvals process for an Inherited Operational Standard from ISB). The Diabetes Summary Core Dataset is subsumed within the DCCR and provides the basis of the National Diabetes Audit which is carried out by the National Clinical Audit Support Programme (NCASP). This dataset was approved as an inherited standard by the ISB in March 2003. Other datasets, generally recommended for operational use at local level are outside *HH*’s scope—these include the Diabetes User Dataset (subsumed within the

DCCR), the Paediatric and Adolescent Dataset and the Diabetes Foot Care dataset. The dataset terms represent a standardised, domain-driven vocabulary shared between end-users and developers. Such terms can form the beginnings of a ‘ubiquitous language’ to build composite service software that more closely follows logical domain models. This is particularly valuable in high-complexity areas such as information governance, individualised risk analysis and longitudinal (intervention outcomes) analysis. Note how the ‘dovetailing’ of data sets map out the domains to provide integrated coverage with common data (much of it from primary care) being shared across domains. The figure is colourised and modified from a UK NHS original emphasising the relationship of the DCCR and *HH* trends monitoring services: Original Copyright © 2007, re-used with the permission of The Information Centre. All rights reserved

depends on resource integrity and consistency of semantics and interfaces plus other constraints regarding legitimacy of access (i.e. security/privacy/confidentiality/ethics are not compromised). *HH* will also take advantage of technology convergence occurring in ‘sensor data’. These constitute the wide array of modalities by which biological (biochemical, biophysical, physiological) absolute measurements/signals/biomarkers can be more consistently captured from individuals into trend services (timestamped episodically, or on a serial-periodic timeline). In an efficient data gathering system, the sensor principles would be embedded in specialised interoperable devices.

As with interoperable web services, global device standards can be leveraged for real-time data flow across multiple applications (Fig. 4). Such an application platform has the potential to interface thousands of ‘niche’ medical and lifestyle-ambient sensors (e.g. specialised measurement devices or alarms) to supply data into multiple advanced service applications. In this scheme, shown in Fig. 4, personalised data hubs on mobile phones, in the home or clinic serve to ‘join-up’ device inputs in any part of the patient’s path. We are currently examining technical options for *HH* data flow into xHR (individualised health records combining functions of personal trend data and clinical trend data).

14 Role-based portlet specialisations for scalable multi-trends monitoring

The *HH* system accommodates three defined end-user roles ‘clinicians’ (in the current research phase, this broadly includes all clinical team roles), ‘patients’ and ‘researchers’. We currently use *IBM Websphere Portlet Factory* for customisation and development of portlets. A portlet, together with its underlying service, is a mini application and each portal page can present a number of portlets each working independently or together. The portlets developed are *JSR168 compliant* supporting portlet interoperability. In the clinician portal, we replicated a task-oriented approach as specified in the Integrated Care Pathway (ICP) for Diabetes. The ICP describes the essential steps in the care of patients with a specific condition and the expected progress of the patient. An ICP is dependent on information that is timely, reliable, secure and specific for a given patient.

The use of biomedical sensor devices (Fig. 4) enables continuous monitoring outside of the hospital setting. *HH* currently uses wireless blood pressure cuff and glucose meter inputs. A point of care patient ID device is used (currently a fingerprint recognition device in conjunction with a mobile hub although this could be a purpose-

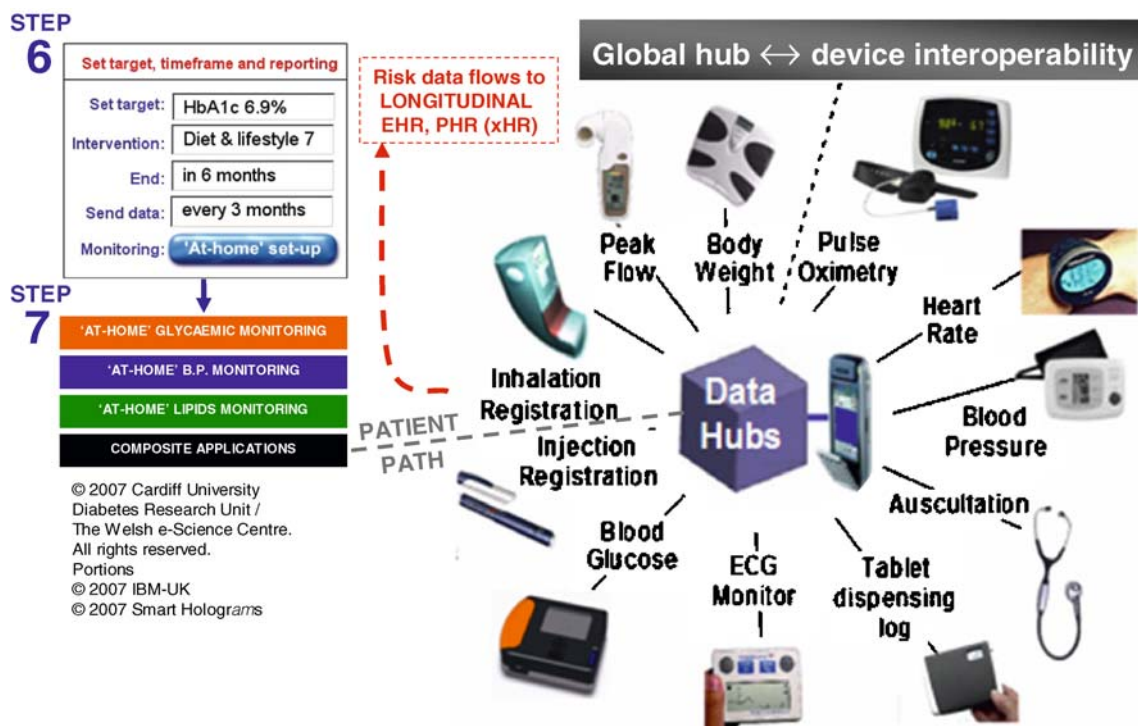


Fig. 4 Device interoperability standards leveraged for real-time data flow across multiple applications in *HH*. The figure illustrates the array of ‘modalities’ by which biological measurements/signals/biomarkers (biochemical, biophysical, physiological) can be consistently captured from individuals into trend services (timestamped

episodically, or on a serial-periodic timeline). In an efficient data gathering system, the sensor principles would be embedded in specialised devices. As with interoperable web services, global device standards can be leveraged for real-time data flow across multiple applications

designed smartcard) to permit multiple patient-per-household usage. We have considered scalability issues in depth—scalability is a principal *raison d’être* for adopting a web services/HealthGrid approach. Figure 5 shows a scheme describing multi-site scalability of the *HH* system (upper). As seen in the figure, each clinical site has its own data centre, and data gathering procedures are made more streamlined and efficient through a standardised workflow called ‘Lean Clinic’ (not described in this paper). Other computational services (yellow boxes) feed information to local and national dashboards (green boxes).

Database schemas have been developed based on the use of DCCR at each participating site (see Fig. 4). The portlets have been developed through a combination of JSP, HTML and *Flash*. *Adobe Flash* forms technology has been used as one alternative to create a dynamic user entry form and for data binding between the form and the Web Services. We have also made use of an open source platform (*Laszlo*, ref. [17]) permitting efficient build/deploy of rich, interactive web applications. *Laszlo* code (XML/JavaScript) when compiled can also be deployed as a *Flash* program. To date, all of the graphical representations in the portal have

been built using *Laszlo*. The graphs provide for a timeline-based layered approach to a key problem: simultaneous trend analysis for evaluating outcomes in patient-centred health monitoring services.

15 Presenting detailed contextual healthcare information in uncluttered and intuitive ways

A major requirement requiring further research is how such a variety of detailed information can be presented in an uncluttered and intuitive manner, with high integrity and ‘appropriateness to role’. For the clinical enduser, multiple trends will likely be visualised in concise ‘at-a-glance’ views that report the patient’s status while making detailed in-context information available by drill-down ‘behind’ timelines and screen objects. In *HH*, detailed contextual information can be obtained by intuitive interaction with graphical elements. For example a time series might exhibit ‘timeline flags’ denoting patient events/activities of relevance to the management of the condition. This approach is drawn from the *Chronicle* metaphor described

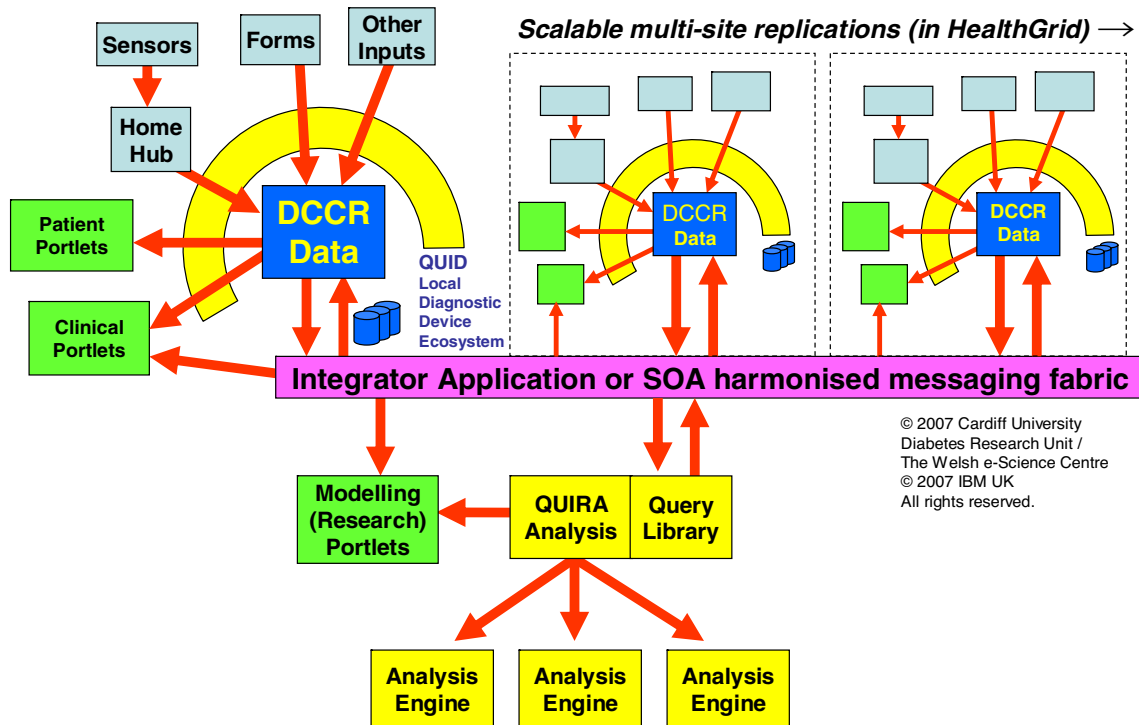


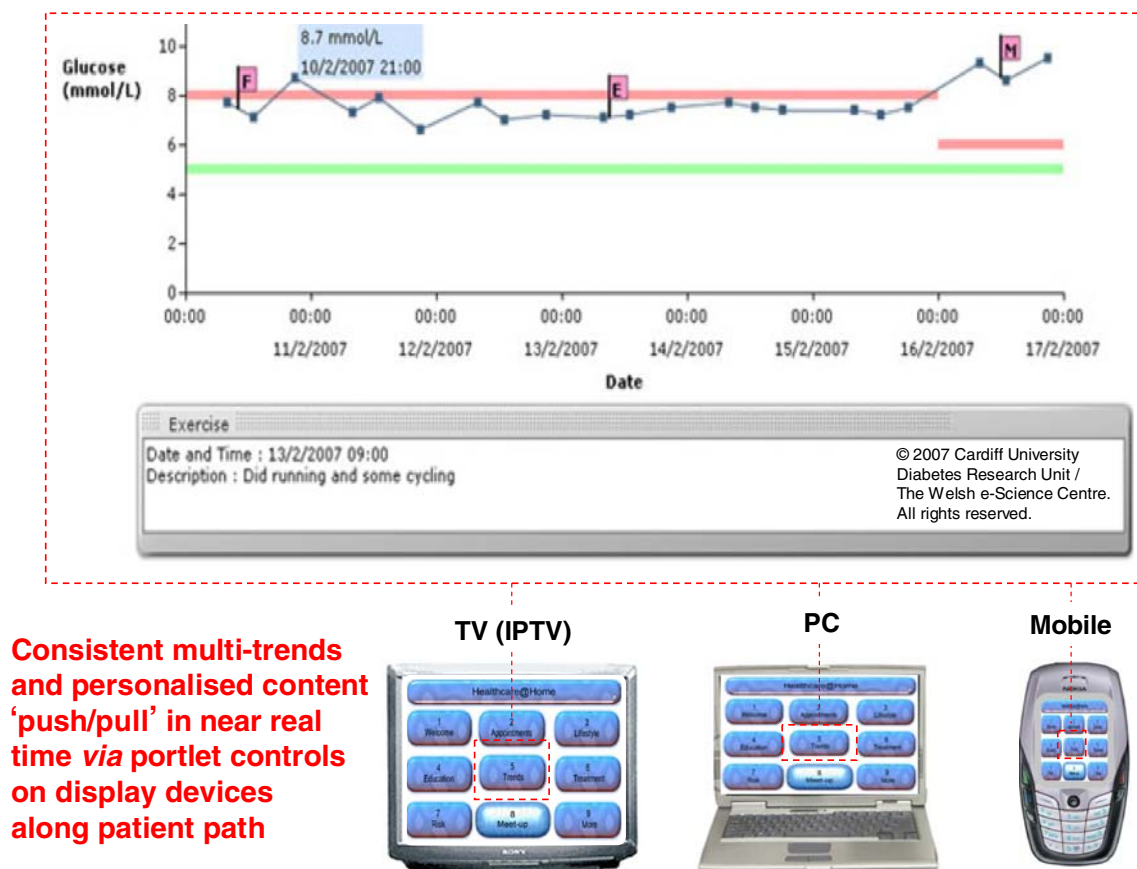
Fig. 5 Scheme for multi-site scalability of the *HH* system. A process server allows data management processes to be deployed and invoked on the server, allowing new processes to be added for additional capability. The process server operates primarily as a workflow enactment engine. The initial request to the process server is generated through Web Services. While the portal replicates the ICP graphically, the process workflow engine captures the business logic. As seen in

the figure, each clinical site has its own data centre, and data gathering procedures are made more streamlined and efficient through a standardised workflow called ‘Lean Clinic’ (not described in this paper). Other computational services (yellow boxes) feed information to local and national dashboards (green boxes). Database schemas have been developed based on the use of DCCR at each participating site (see Fig. 4)

within CLEF (the Clinical e-Science Framework project, see [12, 21]) where data is presented in a textual as well as visual format. In the visual format, timeline-based approaches can be used with multiple ‘stackable, switchable’ layers. Inspired by CLEF, Figure 6 shows a currently-implemented trend service built in *Laszlo* as described above. The flexibility of the *HH* visualisation tools permit patient path reporting over a period of time; this type of functionality can be achieved by federating databases through technologies such as IBM’s Information Integrator (II) or *OGSA-DAI* (Open Grid Services Architecture – Data Access and Integration [30]). We have also carried out an implementation using both II and *OGSA-DAI*. *OGSA-DAI* is a middleware product that allows data resources, such as relational or XML databases, to be accessed *via* web services.

16 Predictive capabilities from quantitative longitudinal data streams

Subject to rigorous ethical information governance and patient control on data reuse, *HH* has significant applications in disease early detection and prevention. Predictive trends services are non-trivial in operation since they need to be ‘scalable’ to cope with potentially massive volumes of incoming time series data. Predictive trends services (e.g. developed for pre-defined pattern matching, specified outcome risk computation and risk stratification can implement outcome risk equations as specialised web services. Authorised clinical roles needing to ‘case find’ for prioritised intervention actions can query a registry of such services and run the risk equation against local case data. This establishes an ‘extensibility’ principle - new



Consistent multi-trends and personalised content ‘push/pull’ in near real time *via* portlet controls on display devices along patient path

Fig. 6 A currently-implemented longitudinal trend and event commenting service built in an open source platform. *Laszlo* code (XML/JavaScript) can be compiled for deployment as a *Flash* program. Timeline-based layers enable simultaneous multi-trends analysis as part of outcomes-based monitoring services through personalised portals. The flexibility permits intuitive, uncluttered ‘patient path data’ reporting over a period of time; this type of functionality can be achieved through database federation technologies (see text). The lower portion of the figure shows options for display of trends via portlets hosted on different types of display device along the patient

path (see *step 7* on Fig. 1). The scope of *HH* data display is currently TV (to an IPTV client), PC (web pages) and mobile devices. Active portlet data ‘push’ and ‘pull’ to/from SOA services and HealthGrids can be combined with consistent, high quality data accession interfaces. For example, keying a specified number—on either a TV remote control or mobile phone—could efficiently access specific types of trend data ‘on the fly’. Data can be delivered to the enduser in high quality embedded *Flash* elements appropriate for that target display device with AJAX extensions for optimising bandwidth

equations can be ‘wrapped’ as they are discovered from conventional outcomes-based research. We are currently thinking through how a dedicated *research outcomes portal* (not shown) can be best designed to supply anonymised data to authorised statisticians, clinical researchers and healthcare policymakers for service-based trends analysis (subject to ethical constraints on re-use of data being satisfied). There is also value in the automating manual aggregation (or transformation) of service or clinical facility data (through federated interoperable mechanisms presently seen in commercial business intelligence tools). These research functions would be especially useful for evaluating the longitudinal impact of research or health policy interventions across whole populations.

17 The patient portal as a key vehicle for patient’s gaining insight into their conditions

The *HH patient portal* component (not shown) has been designed and implemented to (1) address presentation of information in an accessible manner when ‘at home and mobile’ (2) to enable the patients to follow prescribed treatment and guide them when necessary (3) to enable patients to provide additional information on their daily activities that may have a bearing on their condition (4) to make the portal accessible from a variety of devices including mobile phones, PCs and the most common display device in the home—the TV. Data can be visualised as graphs, logs of daily activities, and document pointers (or references) to materials recommended for reading by the clinical team (e.g. doctors, nurses, pharmacists, dietitians etc.). ‘Patient portals’ can also host structured questionnaires that can be used, with ethical guidance, to input risk data. The current HH interface has a simple design so it can be accessed from various devices including: mobile phones, internet browsers and IPTV clients.

The overall design purpose from the outset was that a multimedia presentation layer ‘would help patients gain better insight into their conditions’. A major development in the UK NHS directly related to all portal functionalities has been the Common User Interface (CUI, ref. [29]) project, designed to help ensure that critical patient data such as drug information is displayed in a standard way; the CUI (built on.NET with AJAX extensions) gives a standard patient overview, together with prescribing screens that enable clinicians to easily identify the right patient and ensure they receive the correct medications. Tools for meta-tagging clinical notes (e.g. automatic prompting for SNOMED-CT coding) are included as a contribution towards developing an integrated EHR. Other key aspects of this project include standardised date/time formats, consistent UI, availability of a design guide, medications

resource integration (e.g. British National Formulary) and a toolset for independent software vendors to incorporate the CUI into their own application front ends. *HH* will utilize the outputs of the CUI project in due course.

18 Implications of the *HH* ‘e2e’ scalable system for modelling outcomes

The research underlying the *Healthcare@Home* (HH) information system described here has been expressly focused on a purpose of ‘near real time’ ethical risk analysis for disease early detection and prevention. We have argued that to fulfil this system purpose, achievement of universal technical interoperability (e.g. sensors/transducers/patient path hubs/messaging/storage/retrieval systems) is a vital step - but is only the first part of the solution. Complex health service applications place additional fastidious requirements on the design of information services. Issues arise during the design of safe and ethical ‘end-to-end’ information systems where multiple strands of technology converge. Moreover, the deep contextual complexity and sensitivity of health and social care information makes ‘e2e’ information system design (for efficiency and efficacy) especially challenging. The single ‘ecosystem’ approach has a profound challenge in limiting ‘emergent’ complexity from multiple interacting technology strands as they are added to the ecosystem. *HH* thus introduces principles of *complexity constraint* and *domain-driven design* through policy-matched portlet functions (as part of the ‘closed loop’ to synchronize multi-site convergence) making disease prevention best practice policy deployment easier to ‘scale’.

We have also argued that to deliver patient and service benefits consistently without clinical compromise, a rigorous and unambiguous *international* ethical framework on data reuse is required (explicitly related to the deployment of global ‘disease early detection and prevention’ policies). As data can be more readily exchanged and aggregated through adoption of powerful global technical interoperability standards, more ethical safeguards need to be in place to prevent data misuse. Moreover, while time-based analysis is critical for routine workflow scheduling, resource/capacity management and planning, *consistent* performance metrics related to disease monitoring outcomes are required. The same considerations need to be given to ethical data reuse by quantitative modelling approaches using (for example) longitudinal statistical analysis, algorithmic programming, risk modelling and real-world data-driven healthcare system simulations. There is thus an interdisciplinary task to ‘think through’ information system needs.

We argue that a *global* perspective on semantic consistency will give many additional benefits in the long

run. *Healthcare@Home* makes its technology choices on the basis of technology that is ‘fit-to-need of system purpose’, anchored by a *sustainable and consistent technology-independent articulation* of that need in stable, system-wide, semantically-consistent workflows. This model for technology convergence looks first to global standards (if they exist) but is pragmatic enough to realise that standards are both multiple and in flux (i.e. often take many years to stabilize). In the next phase, however, harmonisation initiatives such as CDAr2 (see refs. [8, 11, 18]), IHE (Integrating the Healthcare Enterprise, latterly supporting interoperability through XDS and XDS-I [7, 31]), the CHA (Continua Healthcare Alliance, [24]) and Common User Interface (CUI, [29]) and more will play significant roles in realizing sustainable person-centred disease prevention services.

19 Implications of the *HH* ‘e2e’ scaleable system for quantitative modelling of healthcare outcomes

Moreover, standards-based interoperability initiatives like these could positively impact ‘affordability and effectiveness’ (cost-benefit) of advanced decision support services. We suspect technical interoperability is only the first step in the journey. Semantic inconsistencies that arise from place-to-place and vendor-to-vendor will be a greater barrier to introducing improved (preventative) ways of working across large populations of individuals. Semantic interoperability—ensuring the same concept is represented identically wherever a device or service records or processes information—is *de rigueur*—and is the critical starting point to guarantee patient safety where multiple devices and software solutions need to work cooperatively. It is a non-trivial problem to ‘scale’ such information services safely, robustly and ethically to monitor large numbers of patients at various stages from healthy (pre-disease) to severe complications and is equally applicable to multiple conditions. There is a need to develop many families of open standard data devices if ‘near real time’ risk data is to be generated and recorded with efficiency.

We have enumerated many data trend analysis routines that may be applied in designing scaleable monitoring systems or quantitative modelling. Subject to rigorous testing, these might have potential to be made available (through an SOA service framework) for processing consented data within a unified semantic interoperability policy. The longitudinal (timeline-based) analytical approach (Fig. 6) also has much potential for analysing distributed data (e.g. expressing the healthcare history of a patient), the workflow (the corresponding processes the patient has experienced) and the log (recording an audit trail of meaningful events during the episode of care).

Subject to ethico-legal clinical data governance and security policies, these data may be distributed among several heterogeneous and autonomous information systems under the responsibility of different local healthcare providers.

An holistic process view such as the *HH* ‘core’ services interoperating with SOA-invoked functionalities provides improved opportunities to integrate data recorded in widely different geographical locations. Distributed queries over different database schemas (facilitated through efforts such as OGSA-DAI—see Section 6) provides a basis for undertaking distributed data analysis. Quantitative modelling (within a semantically harmonised framework) can also discover meaningful outcome-associated patterns within a single measured variables or by correlation (i.e. patterns within multiple variables). Discovered risk patterns can also be compiled to executable equations to ‘case-find’ in local clinical registers. We are continuing to research methods for distributed time series analysis to enable ‘scalability’ of these approaches.

Acknowledgements This project has been funded in the Inter Enterprise Computing Theme of the UK Department of Trade and Industry (DTI)-led Technology Programme, for which we are grateful. We also acknowledge gratefully the contribution each of our industrial partners has made to this project: IBM, Zarlink Semiconductor and Smart Holograms. ECC and DRO are grateful for support by the Wales Office of Research and Development for Health and Social Care, Wales Assembly Government.

Glossary of (Jargon) terms in the context of the project for non-specialists

<i>Aggregate analytical framework</i>	A data collection system that has consistent data formats and collection policies to permit combination of data to suit an analysis task.
<i>Algorithmic modelling of outcomes</i>	A method for discovering the complex dependencies of measured variables to a specified healthcare outcome.
<i>Closed loop outcomes analysis</i>	a key design point of <i>Healthcare@Home</i> , where a high quality risk analysis step takes place prior to and following a specific intervention so to better understand effectiveness.
<i>Composite Process Functions</i>	where information services are separated into small chunks which work together to support a business process.
<i>Domain</i>	in the project’s context, the real-world, day-to-day activities and

	concerns of ‘healthcare’ (as contrasted with technical computing concerns).	<i>Logical Domain Model</i>	in the context used here, a model created by domain stakeholders that organises functional requirements as a framework for technical implementation choices.
<i>Domain Knowledge Expert</i>	someone who has deep understanding of knowledge and its structure associated with the <i>domain</i> .		
<i>Domain-Driven Business Process Modelling</i>	A methodology that captures understanding of domain knowledge experts and other roles associated with the domain. The modelling is useful for the construction of information services that serve the needs of the domain.	<i>OGSA-DAI</i>	Open Grid Services Architecture—Data Access and Integration—A project conceived by the UK Database Task Force with the aim to develop middleware assisting with access and integration of data from separate sources via the grid—see http://www.ogsadai.org.uk
<i>Grid Technologies</i>	<i>in general</i> , the use of resources of many separate computers (working in parallel connected by a network that is usually the Internet) to solve large-scale computational problems.	<i>Portal</i>	In the context used here, a Web site that provides a single point of access to specific applications and information (there may be many portals created for different purposes).
<i>HH</i>	Abbreviation for <i>Healthcare@Home</i> , the name of the research phase demonstrator project being described in this paper—see http://www.ehealthnews.eu/content/view/full/837/27/	<i>Risk Monitoring</i>	In the <i>Healthcare@Home</i> context, continuous or discontinuous analysis of measured variables (e.g. blood glucose concentration) that may indicate specified outcome risk as part of a legitimate healthcare ‘trend’ service. Monitoring can be extended anywhere a telecommunications network can support the devices and applications model.
<i>Integrated Care Pathway</i>	<i>in general</i> , a multidisciplinary outline of anticipated care, placed in an appropriate timeframe, to help a patient with a specific condition or set of symptoms move progressively through a clinical experience to positive outcomes.	<i>Scalability</i>	<i>in general</i> , the ability to expand a computing solution to support large numbers of users (i.e. increasing demand) without impacting performance. Scalable services have potential for reducing cost, though this is untested.
<i>Inter-Enterprise Computing</i>	creation and use of a wide range of environments, equipment and infrastructure in a way that allows secure, appropriate sharing of knowledge and powerful computing resources. The sharing can occur over a wide area using a high bandwidth communications infrastructure (sometimes but not exclusively referred to as grids).	<i>Service-Oriented Architecture (SOA)</i>	put simply, a concept for a software infrastructure that defines how collections of distributed information services can interact to communicate and interoperate <i>via</i> agreed standards. Collections of services are so-called ‘loosely-coupled’ to promote re-use of components and when combined, support complex processes.
<i>Interoperable</i>	<i>in general</i> the ability of systems to exchange and use information, the term is often applied to products and systems from multiple vendors that can be used together without modification or development of custom interfaces and tools.	<i>Time Series Analysis</i>	an approach to identifying underlying behaviour from a

timeline-based sequence of observations, e.g. sensor-derived biomedical data. *TSA* may forecast (predict) future values of the time series variable, but this requires that patterns of observed time series data are identified and formally described.

Web Services

in the context used here, *WS* represent simple, self contained applications which perform specific functions. *Web Services* describe a standardized way of integrating Web-based applications using the XML, SOAP, WSDL and UDDI open standards over an Internet protocol backbone.

References

- Albert PS (1999) Longitudinal data analysis (repeated measures) in clinical trials. *Stat Med* 18:1707–1732
- Baird J, Turisco F (2005) Service-oriented architecture: a critical technology component to support consumer-driven health care. *AHIP Cover* 46:74–76, 78 passim
- Benson T (2007) Prevention of errors and user alienation in healthcare IT integration programmes. *Inform Prim Care* 15:1–7
- Bescos C, Schmitt D, Kass J, Garcia-Barbero M, Kantchev P (2005) Interoperability and HealthGRID. *Methods Inf Med* 44:190–192
- Breton V, Blanquer I, Hernandez V, Jacq N, Legre Y, Olive M, Solomonides T (2007) Roadmap for a European healthgrid. *Stud Health Technol Inform* 126:154–163
- Bucur A, Kootstra R, van Leeuwen J, Obbink H (2006) Service-oriented architecture for grid-enabling medical applications. *Stud Health Technol Inform* 120:55–68
- Dogac A, Laleci GB, Aden T, Eichelberg M (2007) Enhancing IHE XDS for federated clinical affinity domain support. *IEEE Trans Inf Technol Biomed* 11:213–221
- Dolin RH, Alschuler L, Boyer S, Beebe C, Behlen FM, Biron PV, Shabo Shvo A (2006) HL7 clinical document architecture, release 2. *J Am Med Inform Assoc* 13:30–39
- Donnelly K (2006) SNOMED-CT: the advanced terminology and coding system for eHealth. *Stud Health Technol Inform* 121:279–290
- Estrella F, Hauer T, McClatchey R, Odeh M, Rogulin D, Solomonides T (2007) Experiences of engineering grid-based medical software. *Int J Med Inform* 76:621–632
- Ferranti JM, Musser RC, Kawamoto K, Hammond WE (2006) The clinical document architecture and the continuity of care record: a critical analysis. *J Am Med Inform Assoc* 13:245–252
- Hallett C (2006) Summarisation and visualisation of e-health data repositories. Part of Clinical e-Science Framework (CLEF) project. UK e-Science All-Hands Meeting 2006, available at: <http://oro.open.ac.uk/5261/01/AHM06.pdf>
- Hartswood M, Jirotko M, Procter R, Slack R, Voss A, Lloyd S (2005) Working IT out in e-Science: experiences of requirements capture in a HealthGrid project. *Stud Health Technol Inform* 112:198–209
- Hernandez V, Blanquer I (2005) The grid as a healthcare provision tool. *Methods Inf Med* 44:144–148
- Herveg J (2007) Does healthgrid present specific risks with regard to data protection? *Stud Health Technol Inform* 126:219–228
- Kawamoto K, Lobach DF (2007) Proposal for fulfilling strategic objectives of the U.S. Roadmap for national action on clinical decision support through a service-oriented architecture leveraging HL7 services. *J Am Med Inform Assoc* 14:146–155
- Laszlo Systems (2007) Source for OpenLaszlo, an open source platform for rich Ajax application development. <http://www.laszlo.com/> last accessed: June 2007
- Muller M, Frankewitsch T, Ganslandt T, Burkle T, Prokosch HU (2004) The Clinical Document Architecture (CDA) enables electronic medical records to wireless mobile computing. *Medinfo* 11:1448–1452
- Nadkarni PM, Miller RA (2007) Service-oriented architecture in medical software: promises and perils. *J Am Med Inform Assoc* 14:244–246
- Olive M, Rahmouni H, Solomonides T (2007) From healthgrid to SHARE: a selective review of projects. *Stud Health Technol Inform* 126:306–313
- Rogers J, Puleston C, Rector A (2006) The CLEF chronicle: transforming patient records into an E-Science resource. Part of Clinical e-Science Framework (CLEF) project. UK e-Science All-Hands Meeting 2006, available at: <http://www.allhands.org.uk/2006/proceedings/papers/688.pdf>
- Steinichen P, Turisco F (2006) Service oriented architecture: adding value to legacy IT systems. *AHIP Cover* 47:54–56, 58–59
- Ure J, Procter R, Martone M, Porteous D et al (2007) Data integration in ehealth: a domain/disease specific roadmap. *Stud Health Technol Inform* 126:144–153
- Various authors (2007) Continua healthcare alliance; see <http://www.continuaalliance.org> Last accessed: June 2007
- Various authors (2005) Diabetes Continuing Care Reference dataset (DCCR); for details see <http://www.ic.nhs.uk> and <http://www.ic.nhs.uk/our-services/improving-data-collection-and-use/datasets> Last accessed: June 2007
- Diabetes NSF (Wales) (2003) Standards <http://www.wales.nhs.uk/sites3/home.cfm?orgid=440> Last accessed June 2007
- Fifth IMA International Conference on Quantitative Modelling in the Management of Healthcare (2007) Institute of Mathematics and its Applications. <http://www.ima.org.uk/Conferences> last accessed June 2007
- HealthGrid annual conferences, e.g. <http://lyon2003.healthgrid.org>, <http://clermont2004.healthgrid.org> <http://oxford2005.healthgrid.org/>, <http://valencia2006.healthgrid.org/> <http://geneva2007.healthgrid.org>, URL's last accessed. May 2007
- The NHS Common User Interface (CUI) Programme. Link accessed: <https://www.cui.nhs.uk/Pages/NHSCCommonUserInterface.aspx>, 2007
- OMII UK (2007) The OGSA-DAI Project—details at <http://www.ogsadai.org.uk/>. Last accessed: June 2007
- Wozak F, Ammenwerth E, Breu M, Penz R, Schabetsberger T, Vogl R, Wurzl M (2006) Medical data GRIDs as approach towards secure cross enterprise document sharing (based on IHE XDS). *Stud Health Technol Inform* 124:377–383
- Evans E (2003) Domain-driven design: tackling complexity in the heart of software. Addison-Wesley