Alchemist Multimodal Workflows for Diabetic Retinopathy Research, Disease Prevention and Investigational Drug Discovery

Adina RIPOSAN [a, b, 1] Ian TAYLOR [c, d] David R. OWENS [e] Omer RANA [c] and Edward C. CONLEY [c, e]

[a] *Department of Applied Informatics, Military Technical Academy*
[b] *Contact Net Ltd, Bucharest, Romania*
[c] *School of Computer Science, Cardiff, UK*
[d] *Center for Computation and Technology, LSU, USA*
[e] *Diabetes Research Unit, School of Medicine, Cardiff University*

Abstract. In this paper we present mechanisms for imaging and spectral data discovery, as applied to the early detection of pathologic mechanisms underlying diabetic retinopathy in research and clinical trial scenarios. We discuss the *Alchemist* framework, built using a generic peer-to-peer architecture, supporting distributed database queries and complex search algorithms based on workflow. The *Alchemist* is a domain-independent search mechanism that can be applied to search and data discovery scenarios in many areas. We illustrate *Alchemist*'s ability to perform complex searches composed as a collection of peer-to-peer overlays, Grid-based services and workflows, e.g. applied to image and spectral data discovery, as applied to the early detection and prevention of retinal disease and investigational drug discovery. The *Alchemist* framework is built on top of decentralised technologies and uses industry standards such as Web services and SOAP for messaging.

Keywords. Biomedical Informatics, Biomedical integration, Biomedical imaging, Spectroscopy, Diabetic Retinopathy, Search engine, Resource discovery, Grid, HealthGrid, P2P, peer-to-peer, WS-RF, WSPeer, Ubiquitous computing

1. Introduction

The majority of Internet search engines are based primarily on crawling the Web to create large index databases which are sorted according to a sophisticated ranking system. An end-user searches these massive databases via a Web site and retrieves potentially hundreds of thousands of results, sorted according to the given ranking system. The *Alchemist* infrastructure described here uses an alternate approach that allows users to *proactively* push information into a decentralised "search database" using standardised

[1] Email: adina.riposan@contactnet.ro, Web: www.contact-net.ro

Web Services interfaces, developed within the business and Grid computing communities, and hosted on a peer-to-peer (P2P) infrastructure.

The *Alchemist* is a project at Cardiff University that is creating a framework to provide a P2P layer for supporting pluggable network discovery and caching overlays coupled with the ability to execute distributed workflows [1]. The framework is being built on an existing middleware system, called WSPeer [2], which provides a SOAP messaging layer (using Web or WS-RF Services) within a P2P network that supports a super-peer topology of rendezvous or advert caching peers to support the scalability of the discovery and access to information on the network as a whole. This system already interfaces with existing Grid middleware (e.g. Globus) through Web Services interfaces but is also able to provide access to such capabilities within a decentralised environment. For example, rather than relying on centralised discovery mechanisms e.g. UDDI or similar, we can search super peers for WSDL files to provide access to the distributed services. In this way, the network can cope with far more transient participants, it can support exchanging roles and is capable of scaling proportionately with the number of peers.

The *Alchemist* infrastructure can be customised to be used in particular fields and particular user-environments and scientific communities, such as the biomedical research environment, astrophysics research, audio research, etc. This paper here presents an analysis of how such technologies may be used in the field of diabetic retinopathy by aggregating different sources of information at different levels in order to achieve more reliable and accurate results for the patients. The next section describes the application of *Alchemist* to data discovery and diabetic retinopathy workflows; we propose two scenarios - one for early detection and prevention of retinal disease, and one for retinal drug discovery. In section 3, we discuss the *Alchemist* framework and how the distributed interactions take place.

2. Workflows for Data Discovery in Diabetic Retinopathy

Here we propose the study of diabetic retinopathy as a focus for customised utilisation of the *Alchemist* multimodal search/workflow system - one that spans both fundamental biomedical research and routine clinical (screening) practices.

The development of *Alchemist* applications relevant to diabetic retinopathy is being guided by expertise within the Diabetic Retinopathy Screening Service for Wales (DRSSW), and the Diabetes Research Unit (DRU, Llandough Hospital) with the School of Optometry and Vision Sciences (SOVS), Cardiff University. Within the DRSSW and DRU a large body of high quality research data has been collected relating factors associated with specific graded outcomes of diabetic retinopathy (DR). Additional data is currently being obtained from primary care settings. Anonymised retinopathy data collected over time with specific ethical permissions for re-use represents a unique research resource. Data can include high-resolution retinal images plus associated quantitative physiological, demographic and other variables of relevance to an individual's risk of disease progression. Expert manual analysis of selected data within the DRU suggests the existence of progression *patterns* (progression risk signatures) that underlie common disease mechanisms (relevant to several disease endpoints, not limited to DR). Enough has been learned from manual inspection of *disease progression patterns* to suggest the feasibility of computer-based progression risk identification. Disease progres-

sion *data discovery* and distributed *pattern matching* across whole populations has the potential to radically improve efficiency in patient selection, intervention effect evaluation, optimal frequency/cost of screening. However, to realise such advanced retinopathy search paradigms in *Alchemist*, significant work applying focused ontologies needs to be done, introducing formal representations, structured domain terminologies (e.g. from SNOMED-CT constrained by CDA) using a defined messaging specification (such as HL7 V3). Such structures are absolutely necessary to permit interoperability of data amongst families of composite applications in the diabetes domain such as *Healthcare@Home* [3].

2.1. Scenario 1: Early Detection and Prevention of Retinal Disease

A key application of data discovery using *Alchemist* would be in early detection and prevention of disease (in later stages of retinal disease, interventions have diminishing beneficial effects). By performing pattern searches at multiple levels and transformations of signal data *Alchemist* has the potential to help discover - in the outwardly disease-free population - disease associative variables. These could take many forms - signal frequency shifts, consistent patterns, inconsistent patterns, image pixel patterns, physiological variable patterns - that are present in those people that on the timeline *succumbed to retinopathy versus those that did not*. The purpose of *Alchemist*, therefore, could be to discover the variable patterns intrinsic to the disease at the observable pathology level (macro), the molecular and cellular level (micro) and at the nano level. Molecular-level disease-intrinsic biomarkers (e.g. biophysical or chemical probes that generate reliable quantitative of qualitative indicators of disease in combination with an instrument sensor technology) would have the most potential impact worldwide if they can be made inexpensive and reliable. Currently, practitioners are limited with observable indirect pathology that is the result and not the cause of retinopathy.

Physiological domain data patterns from the *Healthcare@Home* project [3] could also reveal disease-associative risk from individual primary care case data (e.g. hypertensive or not, treatments people are on, lipids problems, smoking status, obesity status, HbA1c status, type of diabetes, and duration of treatments). *Healthcare@Home* is the Cardiff University-based research demonstrator for diabetes data interoperability, monitoring, evaluation and pathway-based decision support (see Figure 1). *Healthcare@Home* composite applications provide a research model for data aggregation services and patient-centred care that operate on a large scale using standardised clinical datasets. The *Healthcare@Home* project [3] enables both "push" and "pull" based mechanisms to enable data from mobile devices and/or dedicated home-based network servers to one or more analysis engines. We intend to make these and other data available to *Alchemist* through the *Healthcare@Home* applications so that the search and associative processes can work on them at an individual level. The target system will match continuous data collection from service inputs with intervention episode outcome analysis on the life timeline. For this disease prevention measure, DRSSW has systematically collected data related to retinal grading outcomes on individual timelines (albeit the limited physiology data from primary care sources is currently not integrated). We believe this *evaluated outcome-based* approach to disease prevention that takes into account quality data at multiple levels of observation is an ideal test scenario for the metadata multimedia search potential of *Alchemist*.

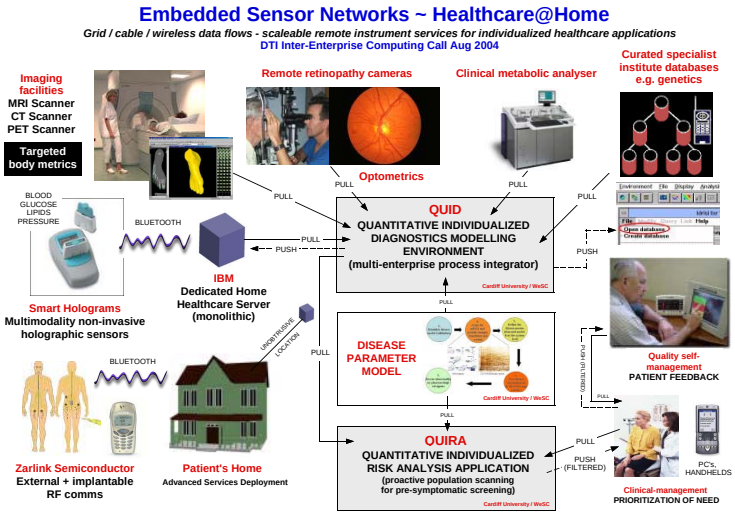

Figure 1. Rich picture illustrating information flow within Healthcare@Home - a research demonstrator for patient-centred healthcare services [3]

2.2. Scenario 2: Retinal Drug Discovery

There is a clear lack of a unified search framework for tracing mechanistic effects of drugs (beneficial and deleterious) on the nervous and vascular cells of the human eye. In its applications to drug discovery and systems biology *Alchemist* has potential to start bridging these wide gaps — e.g. relating drug effect signals observed in relatively simple in vitro cellular test systems with effects in highly complex cellular systems (in man). This comparison framework has applications in reduction of animal testing, especially for drugs whose effects can be tested on the vascular and neural signals measurable in retinal studies. Comparisons could be made between drug efficacy/toxicity effect via (sensor-based) signals on human cellular/molecular dissociated systems versus complex human cell co-culture systems versus in vivo effects recorded within investigational drug settings. A consistent multi-level predictive search framework that could link mechanistic observations at genomic, molecular, cellular, intercellular (tissue), organ and organism levels may better predict failures in drug discovery pipelines. Obtaining good evidence for a drug failure decision early in the pipeline is always preferable to failing a drug later (e.g. when it creates toxicity problems in human subjects).

2.3. Vertical Data Integration in Diabetic Retinopathy

Diabetic retinopathy (DR) is a disease that accounts for c.80-90% of cases of blindness due to diabetes in the UK (c.13% in the overall working population of 19-65 years - Evans Report, HMSO 1991). All persons with diabetes are at risk of developing DR, and its progression to a sight-threatening stage is often not detected. Observable vascular disease in DR - microaneurysms, haemorrhages and cotton-wool spots - have associated *functional* defects that often precede (and therefore can predict) the positional incidence of gross pathology. As reviewed by Antonetti et. al (2006) [4], vascular and neural cell

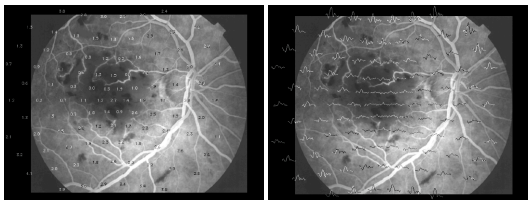

Figure 2. Multifocal electroretinograms (MfERG's) recorded with the visual evoked response imaging system (VERIS) can superimpose localised electrical response amplitudes/latencies and waveforms over digital fundus images and angiograms. For credit see footnote

deficits are inter-dependent in DR, and both are key to understanding the mechanism of vision impairment (see Figure 2[2]).

Moreover, a coordinate marking system should provide for rendering the data comparable between examinations, and correlated within a single time-scale. This could be achieved by an initial point of reference - relating this to the times when the different stages of retinal disease are categorically identified. *At the clinical level*, *retinal fundus grading* is an established method for morphologic assessment of diabetic retinopathy. Grading (i.e. a standardised classification for severity of disease) is generally based on the number, location, and type of *discrete microvascular lesions* in the fundus of the eye. Analysis is generally static for higher lesion densities. At low lesion densities, a chronological sequence of appearance, number, maturation, and disappearance of fundus lesions can be made. Semi-quantitative classifications (based on comparison with standard photographs at various stages of retinopathy) are available but are time consuming and subject to error (i.e. borderline elements can confuse human graders). Detailed lesion mapping / counting and dynamic time series analysis of fundus photographs can be more easily achieved by digital photographic techniques [5].

Although observed microvascular changes are certainly related to retinopathy, *proliferative* retinopathy is better predicted by loss of oscillatory potentials on *electroretinograms* than by methods such as (i) vascular lesions observed on fundus photographs or (ii) capillary non-perfusion visualised by fluorescein angiograms. It is most important to take into consideration the fact that the retina is substantially a vascularised *neural tissue*, not merely a network of blood vessels. *Neural dysfunction* accompanies and precedes visually observable cellular pathology. Numerous reports using electroretinography, dark adaptation, contrast sensitivity, and colour vision tests demonstrate that neuroretinal function is compromised *before* the occurrence of vascular lesions in the retina. From these findings it will be clear that inclusion of data and metadata trails from examinations *at the functional level* is of high significance for developing systems capable of *early detection* and (through early intervention) *prevention* of disease [4].

2.3.1. Automated Retinal Lesion Detection in Diabetic Retinopathy Workflows

Metadata generation in the context of a retinopathy workflow at the *pathology/clinical level* (e.g. colour retina fundus photography grading) consists of contextual observations made by a professional grader. These observations form part of an examination's work-

[2]Permission to use photograph kindly given by Dr. Erich E. Sutter of Electro-Diagnostic Imaging Inc., Redwood City, California, USA

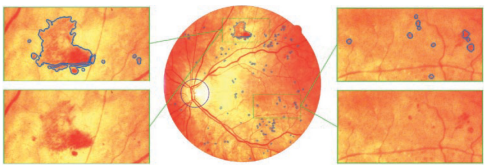

Figure 3. Recognition of retinal lesions by an automated image-analysis algorithm (from Larsen et. al., [5])

flow, e.g. to mark vascular lesions. The grading is then registered on a structured annotation system (see Figure 3) [5].

Figure 3 shows a digitised colour fundus photograph representing non-proliferative diabetic retinopathy with annotation markers outlining red microangiopathic lesion types (haemorrhages and microaneurysms) identified by an automated image-analysis algorithm. Studies have been performed with the objective of developing methods for automated detection of haemorrhages and microaneurysms and to compare this automated lesion detection corresponding to visual identification of diabetic retinopathy. In a key study [5] the feasibility of using automated detection of red lesions in diabetic retinopathy was demonstrated, but the accuracy could be further raised by adjustment of a single user-supplied parameter determining a balance of screening priorities (essentially a piece of metadata). The standard arbiter for automatic lesion detection in the Larsen et. al. [5] study was defined by classifying each patients disease status as having (or not having) diabetic retinopathy (based on overall visual grading of the digitised transparencies). Further assessment in the same patient would indicate whether there has been a change in retinopathy status (based on lesion appearance).

In these and other degenerative pathologies, search functions of *Alchemist* have the potential to reveal if characteristics of graded disease are present. In cases where they are present, *Alchemist* can indicate how they vary over a defined reference timeline.

2.3.2. Systems Biology-based Interpretations of Functional Tests for Retinopathy Status

At the *cellular level*, diabetes alters the function and structural integrity of all retinal cell types. There is evidence for neural function changes in the retina as one of the earliest detectable changes in diabetes. The highly sensitive and complex *trophic* interactions between neuron-neuron and neuron-effector cell signalling in diabetic retinopathy requires a more sophisticated *systems biology* approach to understand. Systems biology approaches (in general) take into account complex interactions of gene, protein, and cell elements to explain system-emergent properties (in the most sophisticated form by biologically realistic modelling of measured outputs of underlying components). For the present time, functional analyses of intact - *in vivo* - systems in the clinic still rely on macroscopic-level techniques. Highly informative molecular-level and cellular-level (microscale) models generally require system-dissociative - *in vitro* - techniques prior to gathering data via (for example) microelectrode and patch-clamp electrophysiology. While these techniques deliver the most detailed mechanistic information, the experimental effects observed do not necessarily scale to the next level up. There remain significant gaps in system understanding essentially due to the dissociative process. The *Alchemist* has some potential to bridge these gaps, but realistically, for the present, data originating from various positional and non-positional clinical optometry techniques have much potential for new search paradigms e.g. multifocal electroretinography

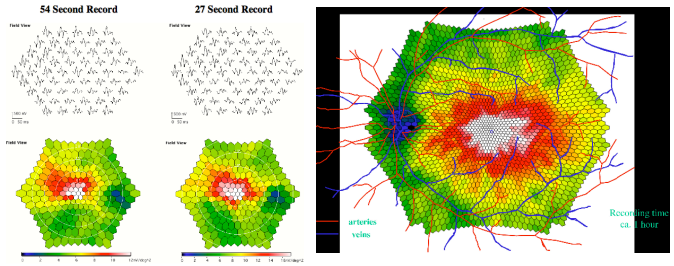

Figure 4. Left: Examples of retinal functional response mapping by electrophysiology in mfERG. Right: Overlay of responses superimposed with positions of arteries and veins. For credit see footnote

(mfERG), multifocal oscillatory potentials (central/peripheral visual field), microperimetry, visual acuity, contrast sensitivity measurements and macula-restricted colour vision.

Investigative ophthalmic techniques as potential data sources for *Alchemist*: *Short Wavelength Automated Perimetry (SWAP)* can detect progressive visual field loss prior to conventional (white-on-white, W-W) perimetry. SWAP is applied in a number of conditions: glaucoma, diabetic macular oedema, various neuro-ophthalmic disorders [6].

For study of intervention treatment efficacies on a longitudinal basis (i.e. analysing individual patients data over time) *2- and 3-dimensional maps of retinal function* represent a powerful direct method. Stimulus-record techniques such as multifocal electroretinography (mfERG) that use *light stimulus patterns* directed across the retina in a predefined 2-d matrix (fig. 4 [3]) can have their corresponding electrophysiological responses recorded within that positional matrix.

There are understandable and sometimes valid issues of lack of trust in data generated by automated routines of any kind. This gap inhibits the development of high-productivity machine learning-based approaches. A significant part of the problem is the lack of *validated datasets* that are of sufficient quality to be used in algorithm performance testing. Consistent outcome datasets that are *trusted* could actually be used for testing and performance ranking (validation) of many algorithms submitted by different computational research groups over the internet (e.g. on a competitive basis for the best blind outcome prediction performance). The purposeful generation of large volumes of interoperable, quality-controlled data is an essential evolutionary step for *consensus pattern searches*, *pattern matching* and *artefact recognition*. Ultimately, a trusted data status for computer-based recognition means signals can be taken as representative of the biological system (and/or the effect of a specified system perturbance).

3. The *Alchemist* Framework and Application to Retinopathy

The *Alchemist* is a P2P framework and an associated workflow toolkit (called the *Alchemist* Integrated Toolkit (AIT)) that provides a flexible way of deploying P2P overlays and allows workflows to be distributed across the system at various levels of granularity. The current system is built on existing well-tested technologies, such as WSpeer, P2PS [4] and the Triana workflow environment [7].

[3] Permission to use photograph kindly given by Dr. Erich E. Sutter of Electro-Diagnostic Imaging Inc., Redwood City, California, USA

[4] http://www.trianacode.org/p2ps/

Triana[5] is a well known workflow environment that can integrate within a number of different distributed environments for allowing true heterogeneous computing across different Grids and distributed paradigms. Such an approach has led to a number of different bindings to underlying middleware and therefore a number of possible modes of operation. For example, Triana on the one hand has a full binding to Java GAT interface, capable of invoking tools and services such as Condor, GridFTP, GRAM, etc., whilst on the other has integrated with service-based middleware, such as Web Services, WS-RF, Jxta and P2PS. It also has the capability to dynamically wrap applications remotely behind Web Services interfaces so that existing software can be easily integrated.

WSPeer has been Triana's Web Services toolkit for the past three years and many projects have used this combination to specify their distributed course-grained service workflows [8]. For example, Triana has been in this capacity[6] for: radio astronomy, astrophysical simulations, data mining, biodiversity problems, grid-enabled medical simulations, environmental science. P2PS has been used as the underlying P2P environment during this time and it has been successfully used in many domains, such as gravitational wave analysis , audio processing and distributed music information retrieval (MIR), distributed P2P simulations and e-health.

WSPeer has a binding to P2PS enabling Web Services to be hosted within decentralised environments and the *Alchemist* builds on this to provide dynamic overlays of peers (like super-peers) that are capable of caching application-specific data (i.e. not just the conventional discovery information) such as metadata or even science data. In *Alchemist* such overlays can be created on-the-fly for the application at hand and deployed onto the peers on the network through the use of P2P groups, used to specify security for participating overlay (e.g. secure authentication, altruistic or reward-based participation, etc). This is necessary for some applications that have sensitive data but wish to employ the use a data-caching overlay of untrusted peers for scalability; the usefulness of such an approach has already been demonstrated through simulations of cycle-scavenging applications, such as BOINC [9]. Here, we are using these mechanisms dynamically propagate retinopathy package workflows (stored as jar files) onto the network using an overlay of package repositories. This scenario is illustrated in our existing DART project, shown in figure 5, illustrating how workflow packages can be propagated through a decentralised layer of package repository cachers (created by *Alchemist's dynamic overlay mechanism*). In DART, users can then query the super peers for music recommendations, being more sophisticated than super peers since they not only cache information but are capable of comparing it too. *Alchemist* makes it very easy for these roles to be defined.

Such a scenario can be adapted easily for retinopathy research and Healthcare@Home, where each Healthcare@Home node propagates information to the network of retinopathy super peers, then can aggregate this information with other available information in order to create higher quality results and healthcare predictions. *Alchemist* can contribute to providing innovative search mechanisms; for example, a multi-modality metadata-assisted search may enable a pathologist to link observations at the level of gross retinal pathology (e.g. vessel bleeds), to physiological dysfunction (e.g. lack of retinal responsiveness), to cellular dysfunction (e.g. hyper-secretory dysfunction), to molecular dysfunction (e.g. constitutive receptor activation), to a genetic

[5] http://www.trianacode.org/

[6] See http://www.trianacode.org/collaborations/index.html for a complete list

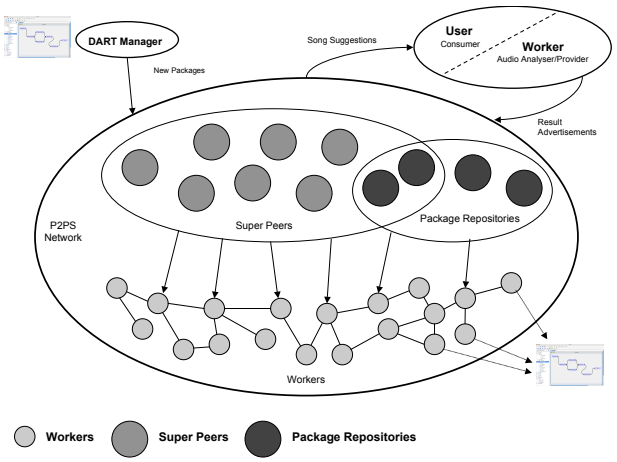

Figure 5. A high level overview of DART and its use of the *Alchemist*, showing the connectivity of its peers.

cause (a specific mutation in the germ-line). In this example, a biomarker that detects the molecular or genetic level could give an unambiguous disease-intrinsic marker that can be applied in future screenings. This benefits the vertical data integration discussed in 2.3, which refers to data representation at multiple levels that might in combination explain origins of pathology (e.g. clinical, functional, morphological, molecular, cellular, genomic). These explanatory combinations may be represented in workflows by analytical algorithms within data discovery workflows or in decision support tools that allow users to apply discovered knowledge. Patterns and signals from combinations of observational modalities (spectral, electrophysiological, retinopathy gradings and annotations/metadata from expert assessments) can help predict what is commonly observed - at the macroscopic level - as a pathological outcome. The high dimensional searches enabled by *Alchemist* also have clear potential here for novel biomarker discovery discussed in 2.3.2 - i.e. identification of signals, technical effects, risk factors or combinations that directly underlie any observation (either in non-exceptional - i.e. normal - functions or in exceptional - i.e. pathological - functions).

By integrating applications, data providers, digital content, and algorithms, the *Alchemist* toolkit enables the simple composition of mixed-media queries for combinational searches to interpret heterogeneous data sets in a logically defined order. Fusing metadata from distinct, yet related, sources into a contextually aware environment allows *Alchemist* to multiplex search results and produce rich metadata that goes beyond what can be traditionally harvested from a single object. At the core, *Alchemist* provides application developers with a framework for specifying complex search algorithms, using a series of logical search steps, within a graphical workflow builder. By building on *Alchemist*'s extensible architecture, developers do not need to write custom software algorithms from scratch, and are able to create complex queries and data fusion techniques in a modular and pluggable fashion.

4. Conclusion

We discuss the *Alchemist* framework, which provides a novel search paradigm based on workflows that can be applied to many application domains for the indexing and search-

ing of digital content. The *Alchemist*, unlike traditional search mechanisms, is based on distributed peers rather than a centralised database. Peers in the network act together to form a decentralised database, where metadata generation and indexing is performed by those peers, avoiding traditional processing bottlenecks. Using the use case scenario for biomedical images and spectral data discovery in diabetic retinopathy, we have described how the *Alchemist* can add value to this community. Annotations added by clinicians or researchers can be used later in complex searches combined with more typical indexing data to produce more detailed and finer grained search results. Annotations by domain experts add value to the stored data, allowing incongruous results to be discarded and non-obvious results to be included. Built on established industry standards such as Web services and SOAP messaging, the *Alchemist* framework and tools are both extensible and interoperable. Reliance on a distributed database and remote processing removes some of the traditional bottlenecks associated with large centralised services or clusters of services.

5. Acknowledgements

We would like to thank our sponsors, PPARC, EU and EPSRC for supporting this *Alchemist* effort through the GridOneD, GridLab and EP/C006291/1 project, respectively. We would also like to thank the researchers and programmers (Eddie Al-Shakarchi, Andrew Harrison, Ian Kelley, Matthew Shields and Ian Wang) at Cardiff for realising these software components. We also thank Healthcare@Home participants at Cardiff University, Smart Holograms, Zarlink Semiconductors, IBM and Llandough Hospital.

References

[1] M. Shields and I. Taylor, "Alchemist: User Driven Searching in Ubiquitous Networks," in *Proceedings of (ADPUC'06) Advanced Data Processing in Ubiquitous Computing*, ACM Press, New York, 2006.

[2] A. Harrison and I. Taylor, "Wspeer - an interface to web service hosting and invocation," in *HIPS Joint Workshop on High-Performance Grid Computing and High-Level Parallel Programming Models*, 2005.

[3] M. Subramanian, A. Shaikhali, O. Rana, A. Hardisty, and E. Conley, "Healthcare@Home: Research Models for Patient-Centred Healthcare Services," in *IEEE 2006 JVA International Symposium on Modern Computing, Sofia, Bulgaria*, IEEE, 2006.

[4] D. A. Antonetti, A. J. Barber, S. K. Bronson, W. M. Freeman, T. W.Gardner, L. S. Jefferson, M. Kester, S. R. Kimball, J. K. Krady, K. F.LaNoue, C. C. Norbury, P. G. Quinn, L. Sandirasegarane, and I. A.Simpson, "Perspectives in Diabetes, Diabetic Retinopathy, Seeing Beyond Glucose-Induced Microvascular Disease," *Diabetes*, vol. 55, pp. 2401–2411, 2006.

[5] M. Larsen, J. Godt, N. Larsen, H. Lund-Andersen, A. K. S. lie, E. Agardh, H. Kalm, M. Grunkin, and D. R. Owens, "Automated Detection of Fundus Photographic Red Lesions in Diabetic Retinopathy," *IOVS, Investigative Ophthalmology and Visual Science*, vol. 44, pp. 761–766, February 2003.

[6] J. M. Wild, "Short wavelength automated perimetry," *Acta Ophthalmologica Scandinavica*, vol. 79, p. 546, 2001.

[7] I. Taylor, "Triana Generations," in *Scientific Workflows and Business workflow standards in e-Science in conjunction with Second IEEE International Conference on e-Science*, (Amsterdam, Netherlands), December 2-4 2006.

[8] I. Taylor, M. Shields, I. Wang, and O. Rana, "Triana Applications within Grid Computing and Peer to Peer Environments," *Journal of Grid Computing*, vol. 1, no. 2, pp. 199–217, 2003.

[9] D. T. Pasquale Cozza, Carlo Mastroianni and I. Taylor, "A Super-Peer Protocol for Multiple Job Submission on a Grid," in *Proceedings of the CoreGRID Workshop on Grid Middleware Workshop in conjunction with Euro-Par, Dresden. To be Published.*, Springer, August 28-29 2006.