Centre for Scalable Data Access in the Oil & Gas Domain Annual Report 2015

— Centre partners and their main Focus —

EXPERIMENT

University of Oslo
Oxford University
Dolphin
Evry
Numascale
Simula
OSIsoft

PROTOTYPE

fluid Operations
IBM
Oracle
Schlumberger Information Solutions
Kadme

PILOT

Statoil
Norwegian University of Science and Technology (NTNU)
Computas

1 Summary

The SIRIUS Centre for Scalable Data Access in the Oil & Gas Domain is a Centre for Research-based Innovation supported by the Research Council of Norway (RCN). The CA for the centre has been signed by all partners and the centre commenced operations on 1st November 2015. SIRIUS consists, at start-up, of a consortium containing Statoil, 10 technology vendors and 4 research institutions. The centre is at present in its start-up period. The formal kick-off and centre opening will be held on 18th-19th May 2016 and activities will ramp-up from summer 2016.

SIRIUS addresses the problems of scalable data access in the Oil & Gas Industry. The centre combines public funding for basic research with funding from its industry partners into an 8 year programme for industrial innovation. It draws together a consortium of leading industrial organisations across the oil & gas value chain, including operators (Statoil), service companies (Schlumberger) and IT companies (Computas, Evry, fluid Operations AG, Dolphin Interconnect Solutions,fluid Operations AG, IBM, Kadme, Numascale, Oracle and OSISoft). In SIRIUS these companies work with researchers from the University of Oslo, NTNU, the University of Oxford and Simula Research Laboratories. The SIRIUS researchers are experts in IT technologies, including high-performance and cloud computing, database technology, semantic technologies and natural language processing. The centre also includes researchers in the area of working practices involving novel technology.

The centre aims to provide the Oil & Gas Industry with better ways to access and use the massive amounts of data that are generated in projects and daily operation. Problems with data access are made more acute by the rise of big data, the internet of things and digitalisation of enterprises. SIRIUS approaches these problems using an interdisciplinary approach, as successful innovation depends on the combination of technologies. The centre is designed to support technological innovation through a portfolio of projects defined by the centre's board. These projects develop basic technology in laboratory projects and then move the technology through prototypes to pilots in industrial applications. The centre's intellectual property model is designed to build a core of open knowledge on which commercial applications can be built.

SIRIUS aims to be an intellectual hub for applied industrial IT in South-Eastern Norway. It has dedicated premises in the Informatics Department at the University of Oslo and collaborates widely, both nationally and internationally. Many of the results generated by the centre will also be of relevance outside the Oil & Gas Industry. The centre is therefore also active in

communicating and applying its results to other areas such as Healthcare, Manufacturing and Public Administration.

Since the centre is in its start-up phase, this annual report serves more as a public prospectus for the centre rather than an account of completed work.

2 Vision/objectives

Over recent years the quantity of available data in the world has been increasing at an exponential rate. The key bottleneck limiting end-user exploitation of this data is the difficulty of accessing *just the relevant* data. As one of the most information-intensive business sectors, the Oil & Gas Industry is affected more than most; for example, Statoil estimates that its annual cost of data access is 250 MNOK in the Exploration unit alone. More broadly, IDC Research estimates that by 2015 the global market for "Big Data technology and services" will reach $16.9 billion, provide 4.4 million jobs, and be growing at an annual rate of 27%.[1]

Recent technological developments offer partial solutions: *High performance computing* and *cloud computing* offer almost unlimited computational resources; novel in-memory and distributed *Database technologies* offer efficient querying of extremely large volumes of data; *Semantic technologies* can integrate heterogeneous data sources and transform data into information; and *Language technologies* support access to unstructured data, including text documents, and its transformation into actionable content.

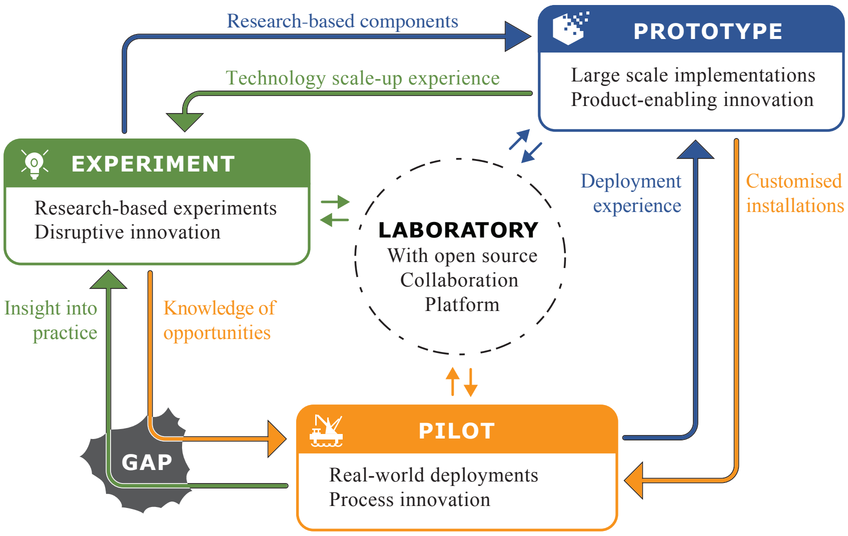

Figure 1: SIRIUS Innovation Cycle.

A comprehensive and robust solution to data access problems requires the integration of multiple technologies; moreover, achieving industry adoption of novel technological solutions will need their integration in professional quality prototypes that can be piloted in real-world deployments. The SIRIUS consortium therefore assembles a team of research and development groups, from both academia and industry, with world-leading expertise in all of the above mentioned technologies, together with major companies in the Oil & Gas sector and leading IT vendors. The consortium consists of partners who cover research and innovation at the level of experiments, prototypes, and pilots. The partners collaborate in the development, deployment, and evaluation of innovative data access technologies, following a value-adding innovation cycle as illustrated in Fig. 1. A shared laboratory and open source collaboration platform provides valuable resources, including real-life software systems and data, that are used for integration, testing, and evaluation at every stage of the cycle; this promotes technology transfer within and between stages, and the transfer of knowledge and expertise between participants.

Crucially, the breadth of complementary expertise within the consortium allows us to close the cycle by bridging the gap that typically exists between research and industry: insights gained from pilots improve the experimental partners' understanding of the application domain, and so inform and guide new research, while collaboration with leading researchers extends piloting partners' knowledge of the opportunities presented by new technologies, and so reduces risks and lowers barriers to adoption. Equally importantly, prototyping partners showcase innovative technologies from the experimental partners by exploiting them in enhanced product prototypes to be evaluated by piloting partners in real-world deployments. This opportunity for collaboration on technology development, with a short feedback loop across the whole value chain, is an unique feature of SIRIUS, and is a major motivation for the participation of industry partners.

The centre builds on experiences from existing EU and RCN projects with established and successful collaborations between academic and industry partners who are eager to join forces and to exploit the opportunities offered by a large-scale centre for research-driven innovation. The structure and organization of the centre is carefully designed not only to cement and reinforce such existing collaborations, but also to encourage new ones that cut across both organizational and technological boundaries, bring together complementary competences and maximize added value. The shared vision of the SIRIUS partners is:

> To accelerate the development and adoption of innovative data access technology in the Oil & Gas industry via broad-based collaboration with a short feedback loop across the whole value chain.

This vision will be continuously elaborated and updated through a bottom-up process, designed to build mutual understanding and ensure genuine commitment from all partners throughout the lifetime of the centre.

SIRIUS exploits the consortium's complementary and wide ranging competences in order to innovate at every level of the technology stack, ranging from cutting edge *computational infrastructure for Big Data*, through *semantically enabled information systems* to *advanced user interfaces* that facilitate new working practices. The consortium's key competences are organised in *research strands*, with research and technology development dynamically structured in collaborative *projects* that exploit expertise from multiple research strands. Projects will be anchored in the centre's objectives, and will have clearly defined resources, workplans and deliverables. To facilitate progression through the innovation cycle, partners from different phases of the cycle will assess new project proposals and evaluate the progress of active projects. The SIRIUS management structure will play a strategic role in ensuring that the overall quality and balance of the project portfolio continues to meet the centre's objectives.

Although SIRIUS is focused on the Oil & Gas industry, the complexity of the challenges found in this domain mean that solutions developed here are relevant in other areas, and the involvement of major IT companies in the consortium will ensure the rapid transfer of SIRIUS innovations to other business sectors. SIRIUS aims to proactively accelerate this process, and to leverage centre funding, by establishing spin-off projects with third party funding, e.g., via the Horizon 2020 program.[2]

Objectives. The vision of SIRIUS is concretised through the Centre's objectives, listed in Table 1. We recognise that it is challenging to create and maintain an environment that encourages and supports collaborative research-based innovation; objectives O1 and O2 address complementary aspects of this process—the transfer of technology and knowledge around the *innovation cycle*—and motivate the *methodology* chosen for the Centre. Objectives O3, O4 and O5 address, respectively, innovation in *working practices*, *information systems* and the *computational infrastructure* for data access; these objectives are reflected in the Centre's *research strands*. Objective O6 addresses the *continuous process of partner alignment and engagement* and Objective O7 addresses *external impact*; these objectives are reflected in the *organisation* of the Centre. The objectives are broken down into subobjectives; where this is natural, key performance indicators for the subobjectives are given in Figure. 2.

Objectives	Subobjectives as measurable targeted outcomes
O1 Accelerate the innovation process for data access in the Oil & Gas domain	**O1.1** Evaluate prototype components in problem owner pilots **O1.2** Evaluate research components on existing products in the development environments of the business partners
O2 Transfer knowledge and expertise via feedback loop in the innovation cycle	**O2.1** Identify constraints imposed by existing tools **O2.2** Identify opportunities for changes in work practices **O2.3** Validate role of prototyping partners' tools
O3 Transform end-user work-practices	**O3.1** Identify barriers (technical, social, cognitive) to uptake **O3.2** Identify mechanisms to assess operational uncertainties
O4 Deliver scalable information systems for accessing disparate data sources	**O4.1** Integrated access to textual/semi-structured/streaming data **O4.2** Scalable access to Big volumes of data (e.g. seismic) **O4.3** Scalable access to real-time streams of sensor data **O4.4** Complex data accessible through end-user interfaces **O4.5** Reduced cost and risk of maintenance and evolution
O5 Deliver scalable, efficient and robust computational environment	**O5.1** Scalable processing and storage of Big volumes of data **O5.2** Processing of real-time streams of sensor data **O5.3** Effectively exploit affordable hardware platforms
O6 Reinforce mutual understanding and shared vision	**O6.1** Elaborate and maintain shared vision **O6.2** Establish new collaborations **O6.3** Track evolving challenges and technologies **O6.4** Achieve equal opportunities and gender equality within SIRIUS
O7 Establish SIRIUS as an internationally recognised Centre of Excellence	**O7.1** Attract additional funding (e.g., from Horizon 2020) **O7.2** Influence future research directions and funding policy **O7.3** Influence society on Big data and data access **O7.4** Establish PhD track combining research and industry skills **O7.5** Influence the international research community

Table 1: SIRIUS objectives and their measurable targeted results.

Business Relevance and Value. From its historical origin as a rough-neck, handcraft based endeavour, the Oil & Gas sector is transforming into an increasingly knowledge- and information-intensive industry. For the Oil & Gas sector in Norway, SIRIUS can potentially create values in the range of several billion NOK by enabling transformation of currently inaccessible data into crucial information for the industry. More available data, improved quality of data, tools to help assess the uncertainty of data, and reduced time for accessing data have direct business value for the Oil & Gas companies, as crucial decisions are almost always made within a short time frame.

- Improved solutions for accessing and combining data would enable improved optimisation of production that could provide a 5% increase in production in fields operated by the company at the Norwegian Continental Shelf;[3] this corresponds to 10,000 barrels per day, which at today's rates amounts to more than $1 million per day.
- Improved tools for data retrieval could free as much as 30% of its experts' time. This time saving alone is estimated to be worth more than 250 million NOK annually;[4] shifting the effort to analysis and decision support is likely to have a much higher value.

Conditions for innovations are not evenly distributed: innovation density is highest in robust industrial clusters. In a Norwegian context, the Oil & Gas cluster is the biggest (in GDP), and has been identified as especially promising.[5] SIRIUS draws on the strength of this cluster by establishing an associated innovation cluster around the massive — and increasing — challenges of information processing in Oil & Gas. Our consortium has been designed to constitute a full-fledged *innovation ecosystem* around information processing in the Norwegian Oil & Gas cluster. The research partners in SIRIUS have extensive experience with innovation projects with industry, very much including the Oil & Gas cluster.

SIRIUS recognises the role of *open innovation*, i.e., open-ended networks of different niche vendors and service and platform providers, in addition to problem owners.[6] A main motivation of the business partners to join the consortium is to be part of this innovation ecosystem, and

KPI title	KPI target value	Objective
# End-User Innovation	≥ 12 pilots demonstrating prototype components to problem owners	O1.1
# Business Innovation	≥ 20 prototypes demonstrating experimental components to business partners	O1.2
# Experiment Learning	≥ 5 revisions to list of known limitations in existing end-user solutions	O2.1
# Pilot Learning	≥ 3 new problem-solving routines within innovation areas	O2.2
# Prototype Learning	≥ 6 new organisational routines validated through use scenarios	O2.3
# Cohesion	≥ 8 workshops with updates to vision document	O6.1
# Industry impact	≥ 20 new working relations to companies or public institutions	O6.2
# SOTA	≥ 8 revisions to list of SIRIUS challenges and objectives	O6.3
# Research Leadership	≥ 12 distinct main supervisors or project managers	O6.4
# Gender Balance	$\geq 40\%$ women at all levels within SIRIUS	O6.4
# External Funding	≥ 40 MNOK annually from other sources than SFI	O7.1
# Influence	≥ 5 (conference keynotes, European level research roadmaps)	O7.2
# Societal Impact	≥ 30 (newspaper articles, media interviews, public committees)	O7.3
# PhDs	≥ 20 started PhD students, of which ≥ 15 finished, and the rest in progress	O7.4
# Scientific Impact	≥ 124 publications in Conferences and Journals with ≥ 20 average field rating and ≥ 15 publications with average field rating over 100, according to MS academic search	O7.5

Figure 2: SIRIUS key performance indicators.

on this basis to innovate also in other sectors where they have market share. Their expectations regarding potential for value creation is detailed in the Centre Agreement.

3 Research plan/strategy

The innovation domains of SIRIUS are provided by the Oil & Gas (O&G) operators at the Norwegian Continental Shelf *and their suppliers*. In the *Exploration* domain, experts investigate selected geographical areas and use relevant information to evaluate the potential for undiscovered hydrocarbons. In the *Operations* domain, experts make decisions about the maintenance and optimisation of production facilities based on a complex collection of historical data and real-time data streams from sensors. Exploration and Operations represent two crucial phases of upstream O&G activities. Taken together they cover O&G operators' key concerns for effective, safe, and sustainable activities.

In both innovation areas, experts have to make business-critical decisions daily, sometimes on very short notice, based on evidence extracted from very large and complex data resources. *Access* to data is necessary, but not sufficient, for good operational decision-making; assessing the degree of *uncertainty* is also essential. The challenge of agreeing on the degree of uncertainty is compounded by the presence of so-called silos, i.e., lack of communication across professional disciplines, geographical regions, and work processes. Differences in software, data models, reference systems, and vocabulary reinforce these silos. Table 2 summarizes the situation.

We next focus on the current state-of-the-art and shortcomings in two key areas targeted by SIRIUS: *information systems for accessing data sources* and the *computational environment for data access.*

Information systems for accessing data sources. Providing *end-user access* to large scale *relational* corporate data stores is the main goal of the ongoing EU-funded project "Optique — Scalable End-user Access to Big Data,"[7], in which several of the SIRIUS partners are in-

Domain	Current situation	Shortcomings targeted by SIRIUS
Oil & Gas in general	• Large collections of data in seismic, log files, RDBMS, text documents • Search and access tools for specific – types of data – user groups – anticipated use • Specialized tools, models and reference systems • Data of variable and/or unknown quality	• No coherence between structured (RDBMS), semi-structured (seismic, logs), and unstructured (reports) information → O4.1 • Difficult to integrate information across silos → O4.1 • Difficult to locate on time all relevant information needed for operational or business decisions → O4, O5.1
Working practices	• Critical success factors proposed • Planned, imposed change models • Struggling to achieve sustainable change • Rationalistic decision-making models • Decision-support tools	• Silo-based work practices → O3.1 • Costly routines for data quality and uncertainty assessment → O3.2 • Lack of focus on end-users, and institutionalization of change efforts → O3.1 • Traditional focus on single tools → O3.1
Exploration	• Relational data models with 1,000s of tables, 10,000s of columns • Databases with overlapping domain maintained by different bodies • Access to relational data via SQL; to documents via text indexing; to seismic/log datasets via meta-data	• Complex structure of data makes it accessible only to IT experts → O4.4 • Large variety of representations even for structured data of same domain → O4.1 • Targeted access to large volume data (e.g. seismic) impeded by limits of storage and processing power → O4.2, O5.1, O5.3
Operations	• Existing data warehouse aggregates data – from many sources – from several silos – including high velocity sensor data • New data sources are added to the data warehouse as needed • Many business processes tied to the data warehouse	• Velocity of real-time data (sensors, 4D seismic) exceeds processing and storage bandwidth → O4.3, O5.2 • Variety of data involved makes addition of new data sources complex and error prone → O4.1 • Changes to data warehouse threaten to break existing applications → O4.5

Table 2: State-of-the-art and shortcomings addressed in the O&G domain.

volved. We summarise the current state-of-the-art and shortcomings based on experiences from Optique; these shortcomings are addressed by Objective O4 of SIRIUS. The state-of-the-art is summarised in Table 3.

Computational environment for data access. HPC (High Performance Computing) is a key component to enabling scalable Big Data access. However, bottlenecks at the communication layer severely limit performance and usability. We summarise the current state-of-the-art and shortcomings in Table 4 below; these shortcomings are addressed by Objective O5 of SIRIUS.

User-led innovations. Grounded in recent, important insights from innovation studies, the SIRIUS methodology is designed to enable *user-led innovations* with early ideas and later validation leaning heavily on users.[8] Based on the innovation cycle (see Fig. 1), the method consists of an iterative, three-stage process with a laboratory at the core. The method is aimed at overcoming the gap that in practice often exists between end-users in the Oil & Gas domain and cutting-edge IT research and development teams.

While the eight-year lifespan of SIRIUS makes it essential to maintain agility and plan our research dynamically as challenges and end-user requirements evolve, all SIRIUS research will conform to the planned methodology, and in particular to the innovation cycle. For this reason, we implement research by means of *laboratory and innovation projects*, as explained below. Projects drive the innovation cycle by combining different strands of research, and by addressing one or more of the Centre's objectives. Each project is clearly defined w.r.t. lifetime, partners,

Area	State-of-the-art	Shortcomings targeted by SIRIUS
Semantic technologies	• Theoretical results on ontology-based data access • Research prototypes available • Prototypical semantics-based query interfaces • Integrated Optique platform under development • Initial ontologies for small parts of domains	• Limited support for temporal and geospatial data → O4.1 • Scalable query answering only with weak ontology languages → O4.1, O4.5 • No support for streaming data → O4.3 • Addresses only relational, not semi-structured or textual data → O4.1 • Query answering typically assumes localisation of data → O4.1, O4.2 • Existing domain models of limited use for end-user data access → O4.4
Language technologies	• Tools for morphological and syntactic analysis of text • Research prototypes for negation and uncertainty analysis • Token-based document clustering • Keyword-based document retrieval • Question-answering systems for specific tasks and domains	• Lack of realistic use cases for downstream applications of LT analysis → O4.1 • Limited support for domain adaptation of existing methods → O4.1, O4.5 • Limited semantic understanding of documents (shallow content) → O4.1 • Question-answering systems for the oil and gas domain → O4.4
Database and information retrieval technologies	• Relational databases well-established in research and industry • Popular open-source database (e.g., MySQL) • Powerful commercial products • Domain-specific solutions (e.g., RDBMS)	• Performance & scalability issues for Big Data processing → O4.2 • Multiple data sources poorly co-exist in data access systems → O4.1 • High management costs for complex data centers → O4.5

Table 3: State-of-the-art and relevant objectives for information systems.

Area	State-of-the-art	Shortcomings targeted by SIRIUS
Ultra low-latency unified infrastructures	• Strong knowledge and experience with low-latency communication products • Powerful cost-effective PCIe based products by SIRIUS partners • Significant past contributions from various international research projects (IST, IFU, ITEA)	• Scalability problems with unified network platforms → O5.1, O5.2 • No support for flexible inter-host device access, migration and hot-add → O5.2, O5.3 • Limited understanding of fault-tolerance for tightly coupled systems → O5.2
Scalable in-memory compute platforms	• Full stack of existing interconnection products from SIRIUS partners • Innovative single image shared memory systems from SIRIUS partners • Expertise providing domain-specific scalable solutions	• Heterogeneous multiprocessor environments not scalable → O5.2, O5.3 • Minimal support for communication protocol bridging at processor level → O5.1 • Large in-memory computing capabilities expensive → O5.3
Customised storage and parallel retrieval	• Popular Open Source solutions available e.g. OpenStack, Hadoop • Powerful hardware and software solutions from SIRIUS partners Oracle and IBM • SIRIUS partners active contributors in Open Source community	• Searching 'raw data' prone to degraded efficiency → O5.1, O5.2 • Bandwidth under-utilization for storage equipment → O5.1, O5.3 • Limited and expensive high performance computing capabilities → O5.3

Table 4: State-of-the-art and relevant objectives for computational environment.

tasks, resources and outcomes in terms of measurable objectives and concrete deliverables; these are set out in a short project proposal (see Figure 5), which must be approved by the Strategy Board (see Sec. 4). To facilitate progression through the innovation cycle, partners from different phases of the cycle will help to assess new project proposals and to evaluate the progress of active projects.

The Laboratory. The laboratory consists of shared facilities and resources, including hardware and software, the centre location, and researchers. The laboratory is funded over the *core budget*, including laboratory-related research activities. The results of laboratory activities will be freely shared among the partners. They will as a rule also be published, while software developed in the laboratory will be licensed as open source, cf. Section 4.

An Infrastructure activity will ensure that the centre has access to computing resources and facilities needed to host the Centre Administration, the dissemination channels and website, the SIRIUS Space and the Laboratory.

The Laboratory is the main mechanism for ensuring collaboration, sharing of Laboratory Background and SIRIUS Project Results. The Laboratory will have dedicated resources to it disposal, including researchers, machines, software, problem libraries and data sets. The following Laboratory activities are planned in 2016:

- Phase 1. Scoping of the laboratory using the results of the RM1 Roadmap project and the contributions defined by the partners. Deadline end of June 2016.
- Phase 2. Installation and implementation of the laboratory resources as agreed. Deadline end of December 2016.

Laboratory Projects and Innovation Projects. All activities in SIRIUS are organized in projects. Projects consisting of laboratory activities only, and using laboratory background only, are called *laboratory projects*. Projects that have activities that are not laboratory activities are called *innovation projects*. Innovation projects will be governed by separate agreements regulating background and ownership to results, and the activities outside the laboratory will be funded by other sources than the core budget, typically funding from industry partners or from sources such as Horizon 2020.

Driving the Innovation Cycle. The innovation cycle combines cutting-edge research, tool development and real-world deployment: motivated by real problems from the innovation domains, research insights are used to underpin the development of innovative tools, and experience with pilot deployments is used as a feedback loop to identify new challenges and gaps that could benefit from new research and technology development. To address the challenges of transforming and transporting prototypes out of the laboratory, new innovations are exposed to the full complexity of real-world work practices in a gradual and controlled manner

In the **experiment stage**, software components are developed, configured, and evaluated in the lab.[9] The software components combine generic, open-source tools and IP-protected software from the technology vendors. The research partners have extensive experience with lab-based data access solutions that will be cross-fertilized with solutions from the experimental partners. The flexibility of a laboratory environment allows extensive testing that is otherwise very difficult.

In the **prototype stage**, SIRIUS research partners collaborate with active product development teams of the prototyping partners. The goal of the activity at this stage is to interface components from the laboratory with products and prototypes of the prototyping partners so as to facilitate integration. Typical tasks will be the support for APIs used in industry and in partners' product suites, and support for industry models and information standards. The prototyping stage will allow a more extensive usability testing, as well as preparing for piloting at the next stage. The treatment of IP will be of critical importance at this stage,

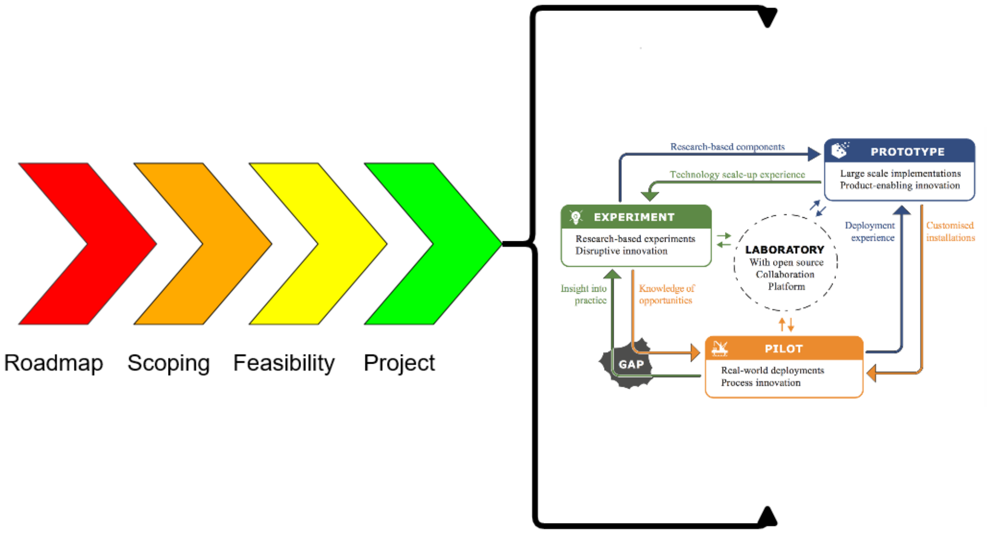

Figure 3: An opportunity becomes a project by passing through four activities.

as the collaboration will rest on a clear border line between open source and the companies' protected IP that has to be established on a case to case basis.

In the **pilot stage**, solutions are installed and deployed at the problem owners alongside their existing tools. The piloted solutions are embedded in the problem owners' existing work practices and evaluated with expertise from the other piloting partners. Crucially, lessons learned are fed back to preceding stages, and the results of prototyping and piloting guide and inform research and experimentation in the laboratory by revealing new challenges and gaps in our understanding.

Premises and Evaluation Method. Table 5 details how evaluation of outcomes from experiments, prototyping, and piloting is done.

Evaluation Site		Evaluation Form
	SIRIUS laboratory	Benchmarks of various types, depending on the type of component. User evaluation with a few selected users.
	On site at vendor	Benchmarking. Usability testing[10]: Selected representative users go through prepared use-scenarios of typical problem situations; user reactions are monitored and evaluated.
	On site at problem owner	Pilots are used alongside existing tools, deployed to a larger community of users and evaluated through case-studies[11] (based on interviews, observations, and logging of use) with expertise from piloting partners.

Table 5: SIRIUS innovation cycle: Evaluation site and evaluation forms.

Managing a Portfolio of Projects 2016 is the year in which SIRIUS establishes a portfolio of on-going research projects. For this reason, we will set up and run a structured program of activities for identifying and capturing project opportunities and developing them into funded, long-term projects and research topics. The pipeline for new projects is shown in the following figure. All SIRIUS projects are governed using this pipeline and we plan to have ongoing activities at all points in the pipeline throughout the life of the centre. An opportunity becomes a project by passing through four activities.

A **Road-map** activity surveys the available technology, technologies and skills of partners and potential international collaborators, and business needs. Its aim is to identify opportunities for research, collaboration and application within the Centre's area of interest. The SIRIUS technology road-map will be prepared by an initial activity in Q2 and Q3 2016. The road-map document will then be revised annually to support the SIRIUS strategy board. Four road-map projects are defined in the centre's work plan. Two of these activities are internal, while the other two ensure coordination with leading international groups: the BYTE project and the Alan Turing Institute.

The broad opportunities identified by the road-map activities will be refined further by a set of **Scoping** activities. In 2016 scoping activities consist of:

1. A general scoping activity that will consist of a planned set of workshops and short-term placements of research workers with partner companies. The aim of this activity is to build a common understanding of the SIRIUS Centre's technological challenges and generate end-user driven proposals for feasibility projects.
2. Scoping activities for each of the six technology strands:
   - Working Practices.
   - Semantic Technologies.
   - Natural Language Technologies
   - Database Technologies.
   - High-Performance Computing.
   - Cloud Computing

A **Feasibility** activity takes a project idea and develops it to the point where a project plan and request for funding can be delivered to a funding organization. Feasibility activities are of limited duration (normally less than 6 months) and generate a clearly defined deliverable: a project plan and proposal for external or internal funding.

All Roadmap, Feasibility and Scoping activities are organized in Work-package 4: Strategy.

4 Organization

Location. The Centre is located at UiO, including the laboratory and Centre Office. Researchers from the other academic partners and the industry participants will be seconded to SIRIUS for longer periods, at least two weeks per quarter and these weeks may be consolidated into longer periods.

Roles and Competencies. The following roles and responsibilities govern SIRIUS:

• *The Chair of the General Assembly* will come from one of the company partners in SIRIUS. The main duties will be to prepare the *General Assembly* meetings together with the Centre Leader, and to chair these meetings.

• *The Centre Leader*, Prof. Arild Waaler, is in charge of the cohesion of the overall strategic direction of SIRIUS with the individual partners' strategies, and thereby also the development of the biennial SIRIUS roadmap and work plan. The Centre Leader acts as secretary to the *General Assembly*, chairs the *Strategy Board*, and participates in the *Operations Board*. Prof. Einar Broch Johnsen is Deputy Centre Leader, in addition to his role as leader of WP4.

• The *Centre Coordinator*, Dr. David Cameron, reports to the Centre Leader and chairs the *Operations Board*. The responsibilities of the Operations Manager are the daily operation of the Centre and its contribution of the projects to the biennial work plan. This includes financial management, maintenance of the risk management plan, and quality control procedures for SIRIUS Project deliverables. The Centre Coordinator acts as secretary to the *Strategy Board*.

• *The Scientific Coordinator*, Prof. Ian Horrocks, provides strategic scientific leadership, aid and advice, monitors the scientific work being carried out within and across projects, and ensures that the work meets the quality standards and overall objectives of the Centre.

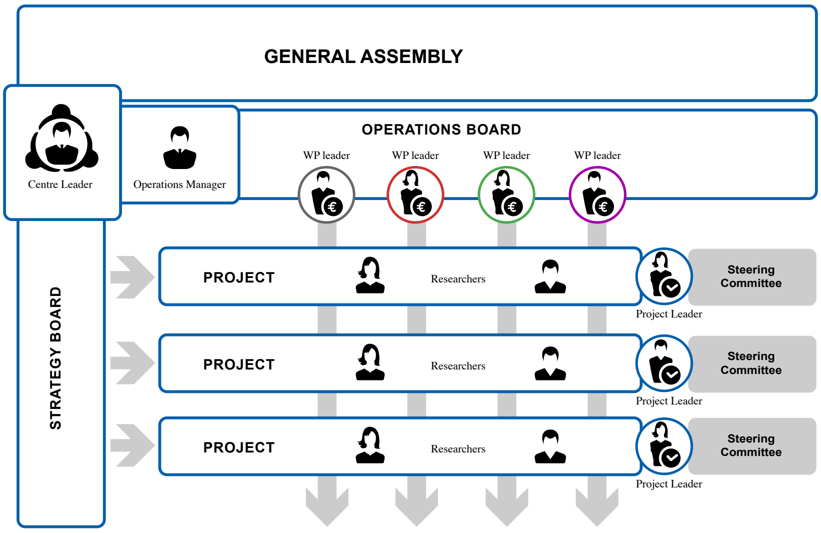

Figure 4: SIRIUS governance structure.

• *The Intellectual Property Manager*, to be nominated by May 2016, is responsible for overseeing the application of the Intellectual Property Rights (IPR) principles of the Centre Agreement and the consortium agreements of individual innovation projects. The IP Manager is appointed by the General Assembly.

• *The Faculty Research Strategist*, Dr. Geir Horn, is appointed by the relevant Faculty at the Host Institution and represents UiO in the Strategy Board as well as SIRIUS towards UiO. The Faculty Research Strategist is responsible for aligning the research strategy of SIRIUS with that of UiO with respect to externally funded research and innovation in Information and Communication Technologies.

• *The Pilot Strategy Coordinator* is responsible for the long term strategy for the pilots ensuring that the pilots serve the interests of the Company Partners as well as providing the necessary input for the research activities in the SIRIUS laboratory.

• *The Mentor and Education Coordinator*, Ass.Prof. Ingrid Chieh Yu, oversees the mentoring framework for the Centre's researchers, and coordinates educational tracks and recruitment of students and Ph.D. candidates.

• *Work Package Managers* ensure that the scientific focus areas of the work packages are covered by SIRIUS's portfolio of projects, and initiate new projects as necessary in order to meet the research objectives of the work packages.

• *Project Managers* lead the SIRIUS projects and manage the day-to-day work and resources assigned to the project.

• *Researchers* work on the SIRIUS projects.

Governance. The governance of SIRIUS, illustrated in Fig. 4, is assured by the following entities:

• *The General Assembly* has one high-level representative from each partner, and makes final decisions. The Assembly will meet virtually as required, and physically on demand. There will be two physical meetings per year focusing respectively on the strategy of SIRIUS and the corresponding rolling biennial work plans. This will foster a mutual understanding of all partners' strategies, harmonise their strategic views on SIRIUS' development, and maximise the exploitation potential of the developed technologies. The Assembly's quorum is 2/3 of the consortium participants, and decisions are normally reached by consensus.

• *The Strategy Board* is responsible for elaborating and implementing the biennial strategic road map for SIRIUS. It is chaired by the Centre Leader and consists of the Scientific Coordinator, the Intellectual Property Manager, the Pilot Strategy Coordinator, the Faculty Research Strategist, and the leader of the strategy and outreach work package (WP4). The Strategy Board is responsible for SIRIUS' project portfolio: The Board prioritises projects according to the strategic roadmap of SIRIUS, formally approves projects, appoints Project Managers, appoints Project Steering Committees, and reports the decisions to the General Assembly.
• *The Operations Board* is responsible for the daily operation of SIRIUS, and for elaborating and implementing the biennial work plans. It is chaired by the Operations Manager and consists of the Centre Leader, the Mentor and Education Coordinator, and the Work Package Managers. The Operations Board is responsible for the financial monitoring of SIRIUS and to ensure timely and correct reporting of the Centre's activities to the Research Council.
• *The Project Steering Committee* consisting of one representative of the Strategy Board, one representative of the Operations Board and one independent representative. They meet every four months to monitor the project progress, and the quality of its deliverables. The Project Manager reports to the Steering Committee.
• *The SIRIUS Office* is responsible for all financial and administrative tasks and the daily operation of the Centre. The Office is led by the *Operations Manager*.

Conflicts will be handled and resolved at the lowest possible level through negotiations among the involved individuals. If the *Project Manager* cannot resolve a conflict within a project, the *Operations Manager* will be notified and a settlement sought in an Operations Board meeting. In the absence of a solution, the Centre Leader will bring the issue up with the General Assembly representatives of the concerned parties. Ultimately, it may be necessary for the General Assembly to settle the dispute definitively.

Technology Transfer. We include general principles for IPR management of SIRIUS Laboratory results in the Centre Agreement and use an open source strategy.

• *Intellectual Property Rights (IPR): General principles.* Any results from a *laboratory activity* will automatically be assigned the status of *laboratory result*. Provisions regarding the *innovation project results* achieved in specific *innovation projects* are stated in their respective *consortium agreements*.

Each consortium participant will have ownership rights to the results produced by that consortium participant, its employees, or suppliers. Access to all laboratory results shall be given free of charge to each consortium participant. Access to innovation project results that are not resulting from *laboratory activities* is regulated in the consortium agreement signed for that specific innovation project.

• *Open Source strategy.* In order to ensure that the generic components developed as part of the research activities and with direct or indirect funding from The Research Council of Norway will be readily exploitable by all SIRIUS partners, it has been decided that *laboratory results* as software will be released under one of the two compatible licenses; the Lesser General Public License (LGPLv3[12]-License) or the Apache Software Foundation License (APACHEv2[13]-license). The consortium participants shall normally choose to release results under the LGPLv3-License, but in the cases where the partners contributing to a joint result cannot agree which license to use, the APACHEv2 license shall be applied.

The two licenses are one-way compatible, meaning that code written under the two licenses can be combined at a component level as long as the derivative code is released under LGPLv3. LGPLv3 is a copyleft license; i.e., modifying LGPLv3 code requires the modified code to be released under LGPLv3. Code with the APACHEv2 license can be changed without any obligation to release the result.[14] Enhancements of mixed code must therefore be released under LGPLv3. Code carrying both licenses may freely be integrated as-is with commercial (closed) software, and there is virtually no difference between the two licenses with respect to use.

Partners SIRIUS draws together a consortium of leading industrial organisations across the oil & gas value chain, including operators (Statoil), service companies (Schlumberger) and IT companies (Computas, Evry, fluid Operations GmbH, Dolphin Interconnect Solutions,fluid Operations AG, IBM, Kadme, Numascale, Oracle and OSISoft). In SIRIUS these companies work with researchers from the University of Oslo, NTNU, the University of Oxford and Simula Research Laboratories.

The planned contributions of each of the partners is listed in Table 6. The centre is currently refining and detailing these contributions through a road-mapping activity.

Partner	Role in SIRIUS
Statoil	Provider of use cases and target for pilots.
Computas	Knowledge of real-world deployment. Facilitating pilots.
Kadme	Provision of leading search and retrieval technology, cross segment E&P data integration and information management domain knowledge.
IBM	Provider of cloud-based software infrastructure. Contributions from leading R&D teams in the US.
Oracle	Big Data technologies, including databases. Contributions from leading R&D teams in the US.
Schlumberger Information Solutions	Leading expertise in software systems for the Oil & Gas sector. Contributions from leading R&D teams.
fluid Operations AG	Contributes leading semantic technology platform and know-how in its deployment.
Evry	Contributes technology and know-how related to hosting of large-scale Oil & Gas data sets.
OSIsoft	Contributes leading technology platform for historians and know-how in its deployment.
Numascale	Big Data technologies, in particular for scalable memories.
Dolphin Interconnect Solutions	Big Data technologies, in particular related to reduced latency.

Table 6: Contributions by SIRIUS company partners.

Cooperation between the centre's partners Initial activities related to cooperation between the centre's partners have been held. These include:

- Collaboration between University of Oslo, University of Oxford, fkuid Operations AG, Computas and Statoil in an application to the EU Horizon 2020 program in the area of process design, operations and maintenance.
- Work with IBM on applications for funding to the EU Horizon 2020 program and the Norwegian Research Council in areas of relevance to the centre.
- Work with Simula, University of Oslo, Dolphin Interconnect, Statoil and Numascale, related to setting up an HPC test environment and defining a research program for approval in autumn 2016.
- Initial meetings between University of Oslo, Schlumberger and Statoil to define research possibilities.
- Cooperation with Computas in participation at the coming Subsea Valley conference in Oslo. Sirius will be represented and profiled at Computas' stand.

The SIRIUS Space at the University of Oslo The SIRIUS centre will be located in a renovated floor of the Informatics Building at the University of Oslo. This renovation has begun in November 2015 and will be completed in Aoril 2016. This SIRIUS Space will provide an environment where:

- All University of Oslo researchers associated with the centre are located together.
- Space is available for visiting researchers from partners and collaborating institutions.
- Facilities and décor will establish a centre identity and provide a pleasant and attractive environment for partners to hold meetings and work when in the Oslo area. Our aim is to make the SIRIUS Space an attractive forum for the Oil and Gas IT industry in South-East Norway.

Objective	Subobjective	Work Package	Strands
O1	O1.1	WP1	S2,S3
	O1.2	WP2	S4
O2	O2.1	WP2	S4
	O2.2	WP1	S1
	O2.3	WP2	S4
O3	O3.1	WP1	S1
	O3.2	WP1	S1
O4	O4.1	WP3	S5,S6
	O4.2	WP3	S7,S9
	O4.3	WP3	S5,S9
	O4.4	WP3	S1,S5,S7
	O4.5	WP3	S5,S7
O5	O5.1	WP3	S7,S8,S9
	O5.2	WP3	S5,S8,S9
	O5.3	WP3	S8

Table 8: Work package objectives.

Partner	S1	S2	S3	S4	S5	S6	S7	S8	S9
Statoil		×	×						
Computas	×	×	×						
Kadme		×	×	×		×			×
IBM		×	×	×	×	×		×	×
Oracle				×			×	×	
Schl. Inf. Sol.		×		×					
fluid Operations				×	×	×			×
Evry		×						×	×
OSIsoft			×	×			×		
Numascale				×				×	
Dolphin Interc.				×				×	×
UiO		×	×		×	×	×		×
NTNU	×		×				×		
Simula Res. Lab.		×	×	×				×	×
Oxford					×		×	×	

Table 9: Strands by partner.

5 Scientific activities and results

	Name	Activities
WP1	Pilots	Strands S1, S2, S3
WP2	Prototypes	Strand S4
WP3	Experiments	Strands S5, S6, S7, S8, S9
WP4	Strategy	Strategy, roadmap, management, quality, risk, communication, exploitation, continuation

Table 7: Work packages.

Activities in SIRIUS will be thematically collected into four *work packages* (WPs), as shown in Table 7. The main purpose of the WPs is to control the balancing of resources between activities. The research and technology development work within SIRIUS is organised in work packages WP1–WP3, which reflect the stages of the SIRIUS Innovation Cycle (Fig. 1) and the SIRIUS methodology. Activities within these work packages are organised in laboratory and innovation projects as described in Section 3. The managerial work package, WP4, is concerned with the RTD strategy and roadmap of SIRIUS, and the day-to-day follow-up management, including quality and risk management, and the communication, exploitation, and continuation activities of SIRIUS. These will be based on regularly revised strategy documents, and address Objective O7: public media activities, spin-off companies, the "Partner Program", and strategic initiatives for additional funding from, e.g., Horizon 2020.

Breakdown of Objectives on Centre Activities. Table 1 presents a break down of the SIRIUS objectives in terms of their measurable targeted subobjectives. These subobjectives can be mapped to the work packages of SIRIUS in a direct way, as shown in Table 8.

- **Work package 1: Pilots** is responsible for SIRIUS objectives related to the evaluation of prototype outcomes for pilots (O1.1) and their integration in working practices (O2.2). Work package 1 also covers SIRIUS objective O3 as this work package addresses uptake barriers (O3.1) and operational uncertainties (O3.2).
- **Work package 2: Prototypes** is responsible for SIRIUS objectives related to the evaluation of experiment outcomes for prototypes (O1.2). Work package 2 also addresses the existing technology restrictions (O2.1) and the industrial state-of-the-art (O2.3).
- **Work package 3: Experiments** is responsible for SIRIUS objectives on scalable information systems (O4) and a scalable computational environment (O5) of SIRIUS.
- **Work package 4: Strategy** is responsible for organising activities addressing objectives O6 and O7 from Table 1.

Organisation of Work Packages. Each of the scientific work packages is organized internally in *thematic research strands* which capture the key competences of the Centre. A research strand brings together experts with relevant competence from different members of the consortium. It is important to observe that the strands are *orthogonal* to innovation projects: projects combine tasks from different strands to achieve objectives and produce deliverables, and thus generate coherence between strands; strands cut across projects to achieve synergies between different RTD tasks. The breakdown of the R&D activities in each work package addressing the different subobjectives, is detailed per strand in Table 8. Observe that several objectives will be realized through synergies between R&D activities in different strands, which is characteristic of research triggered by innovation. Table 9 gives the distribution of the interests of SIRIUS partners over strands, showing that every strand is covered by at least one academic partner and several company partners. Table 6 describes the roles and contributions of SIRIUS company partners, covering both the know-how and technologies they bring into the Centre.

The Work Package Strands. We present the research strands grouped according to the work package structure.

WP1: Pilots. These research strands are related to facilitating the pilots at the problem owners. We foresee research related to challenges posed by the complex interplay between technology and working practices, in addition to the considerable RTD challenges of integrating software components from the laboratory and prototypes into the pilots.

Strand S1: Working practices. Existing work practices exhibit considerable inertia. They are maintained through cognitive (habits), social (routines) and technological (tools) means, and typically resist imposed change efforts. The end-users within the two SIRIUS innovation domains are highly competent knowledge workers that exercise substantial autonomy and discretion. Change efforts through new tools gain legitimacy by grounding them in the needs and demands of the users. Identifying users' needs draws on the expertise and experiences of the piloting partners and invariably takes the form of uncovering (by participant observation, interviews and logging of use) rather than merely 'eliciting' requirements. User requirements are more often than not diverging or conflicting, severely compounding the challenge of uncovering them. The two innovation domains of SIRIUS comprise literally thousands of end-users and numerous sub-groups, many with competing requirements.

An organizational implementation process is exactly that, a *process*. It involves continuously responding to new user demands, facilitating uptake (information, training, support) and resolving barriers arising across silos.

Strand S2: The Exploration pilot. A defining tenet of the Exploration pilot is the uncertainty of data. Exploration is organised as a stage-gated funnel. Starting from highly uncertain prospects, geological interpretations are gradually elaborated by analyses, complementary data sources and collective deliberations (e.g., peer reviewing). This strand will pin-point which, where, and when SIRIUS prototypes are introduced to Exploration experts.

Strand S3: The Operations pilot. In the Operations pilot, a crucial tenet is the extreme variety and velocity (e.g., live streams) of data, data that serves a heterogeneous and distributed community of users. This strand will target selected user communities in Operations to accommodate and fit SIRIUS prototypes to their work practices.

WP2: Prototypes. This research strand is related to technologies that prototyping partners bring into the Centre with the aim of integrating components developed at the experimental stage. These strands focus on lifting software components in an experimental stage from the laboratory to the prototype stage by interfacing these components with products and prototypes of the prototyping partners.

Strand S4: Prototyping. SIRIUS prototyping partners own a portfolio of IT solutions which are well established in the Oil & Gas sector. Products in this portfolio have a significant potential for improvement through the integration of semantic technologies. Tasks contributing to prototyping include

- enabling access to new information sources via semantics-based layer
- using semantic technologies as integration middleware
- developing semantic solutions to reduce installation and maintenance cost
- improving deployment models and scalability through cloud storage and processing

WP3: Experiments. These research strands are mainly driven by the academic partners. They provide the basis of fundamental research and the laboratory implementations that will fuel the prototyping and piloting activities.

Strand S5: Semantic Technologies. This strand addresses technologies that are based on a combination of loosely structured data (e.g., RDF) and conceptual models (ontologies) with a precise and machine processable semantics. In this research strand, we leverage that principle to facilitate both the seamless and robust integration of heterogeneous data sources, and communication between domain experts and information systems. The Optique project is positioned in this area. Typical tasks for this strand include: the extension of Optique technologies to temporal and geospatial data; the combination of relational with semi-structured and textual data; the definition of refined domain models; and semantics-based user interfaces for data access.

Strand S6: Natural Language Technologies. The Language Technology strand aims at providing linguistically sophisticated analysis of textual documents for use in semantically enriched document retrieval. Typical tasks for this strand include: document processing, including automated syntactic and semantic analysis of documents; document clustering, based in part on information provided by processing; document enrichment, including merging documents with structured sources, and extraction of meta-data compatible with a specified ontology; and document retrieval, including experiments with federated search against textual and structural data, and query expansion using semantic categories.

Strand S7: Database Technology. Traditional RDBMS architectures often prove insufficient to cope with the O&G industry's processing and storage requirements, while also being poorly suited to modern cloud processing facilities, and increasingly costly to manage. These shortcomings are addressed in this strand via tasks such as: Big Data extensions (partitioning/parallelisation) to RDBMS; 'always-on' technologies (management, maintenance, and hardware scaling without downtimes); and integration of relational data with RDF, XML and JSON-like data.

Strand S8: High-Performance Computing. This strand addresses technologies to improve computational performance for the applications that require significant QoS guarantees and low latency communication. SIRIUS partners will continue to work on already-solid research foundations to provide solutions enabling direct storage data transfers and GPU data flows to improve application efficiency. Typical tasks for this strand include: innovative ultra low-latency unified infrastructure; scalable in-memory compute platforms; and cost-effective solutions for Big Data domains.

Strand S9: Cloud Computing. The Cloud Computing strand comprises activities, competences and technologies with the aim of providing both new and improved cloud services to the problem owners, as well as to improve the efficiency and flexibility of the cloud architecture itself, implemented by SIRIUS partners. The corresponding research and innovation will typically include tasks like: parallelisation and scaling strategies for data access; improving data

Project outline: Cognitive Advisor for Upstream O&G Activities.	
Access to a wide variety of data is required in Oil & Gas Exploration and Operations, including both public (e.g. NPD) and propriety datasets. This project proposes to exploit HPC and semantic technologies in order to adapt cognitive computing technologies to a variety of such datasets, and create a cognitive advisor solution. The new technology will be tested in the SIRIUS Data Access Laboratory and then piloted.	**Objectives**: O1.1, O4.4, O5.1, ... **Deliverables**: List of deliverables with specified times **Project leader**: P_{pl} **Project steering committee**: P_1, P_2, P_3 **Partners:** P_{pl}, P_2, P_3, P_4, P_5 **WPs:** WP1 (x PM), WP2 (y PM), WP3 (z PM) **IPR and ownership of results:** ... **Total duration:** M12–M48 (36 months)

Figure 5: Sketch of a SIRIUS project cross-cutting strands, based on a proposal from P_{pl}.

RM1: Project Expectations, Scoping and Partner Commitment	
Starting date and duration:	2016-01 for 5 months
Type of Research:	Industrial
Partners involved:	All partners
Background/Rationale: During the inception phase of the centre, each partner will need to define their ambitions and commitment to the Centre.	
Objectives: Establish common perception of partner interests and define the scope and content of each partner's contribution to the Centre.	
Activities current year: Each participant in the project will prepare a position paper that describes: • Their perception of the research state of the art in the area covered by the centre. • Their expectations from the projects run in the centre. • The ways in which they intend to contribute to achieving the aims of the centre. These position papers will be used by the Strategy Board to build more detail into the centre work plan and priorities. Each partner's formal commitment of funding and in-kind contribution (Laboratory Background and in-kind effort) will be made in Appendix 2 and Appendix 3 to the Centre Agreement. These documents are due at the end of March and will be approved at the General Assembly of the Centre in May 2016.	

RM2: Technology Roadmap and Gap Analysis	
Starting date and duration:	2016-04 for 5 months
Type of Research:	Industrial
Partners involved:	IBM, Oracle, Numascale, fluid Operations, Evry, Dolphin, Kadme, Computas, OsiSoft, Schlumberger
Background/Rationale: Create a common understanding of the technologies covered by the centre and conduct a gap analysis so that the centre's project portfolio is optimized.	
Objectives: Prepare a report that presents a technology roadmap and gap analysis. This document will be used by the Strategy Board to initiate Scoping and Feasibility projects from Q3 2016.	
Activities current year: A team of technology vendors will work with the project manager to prepare a report covering the objectives described above.	

access on-site through distributed cloud architectures, making for efficient data access out in the field; off-site big data analytics accessed through cloud services; cloud federation to seamlessly utilize resources from clouds with different characteristics, as well as for flexible scaling between private and public clouds; utilization of High-Performance Computing technologies as a foundation for improved cloud architectures more efficiently supporting big data and big data analytics.

An example of a SIRIUS project. The short project description in Fig. 5 illustrates how a project in SIRIUS integrates technologies and/or expertise from different strands.

Status of Roadmap Activities The objectives and plans for the two initiated road-map activities, RM1 and RM2, are shown in the following tables.

Status of Scoping Activities Seven scoping activities are planned in 2016. The first of these is a general scooping activity that will support a program of workshops with partners and short term placements of research personnel in a partner's organization, as part of their research work. The purpose of this structured contact is to:

F1: Next-Generation Operations	
Starting date and duration:	2016-01 for 6 months
Type of Research:	Industrial
Partners involved:	Statoil, UiO, Computas, fluid Operations AG, University of Oxford
Background/Rationale: The Optique project identified opportunities for application and piloting of the developed methods in process operations. These fundamental and industrial challenges will be addressed by a follow-up piloting project which will apply for funding from the European Union in April 2016.	
Objectives: Define a project plan and establish a consortium – containing both SIRIUS partners and external parties – that will work on a long term piloting project on application of Ontology-based Data Access to operational problems in the process industry.	
Activities current year: Establishment of a consortium to apply for EU funding at the Horizon 2020 call in April 2016.	

F2: Natural Language and Scalable Data	
Starting date and duration:	2016-01 for 6 months
Type of Research:	Fundamental
Partners involved:	IBM, UiO
Background/Rationale: Work on Scalable Data Access thus far has concentrated on structured and streaming data, while text and unstructured data has not been addressed. It is therefore worthwhile to explore how Scalable Data Access technologies can be combined with Natural Language methods to allow access and alignment of information across structured and unstructured data sets. This approach can be applied in multiple industries.	
Objectives: The SIRIUS centre is participating in an application to the RCN IKTPLUSS Fyrtårn programme in health. SIRIUS researchers will contribute to this program with research and skills in Ontology-based Data Access and Natural Language Processing. IBM is also participating in this proposal.	
Activities current year: Participation in preparation of funding proposal.	

- Build relationships between research workers and operational personnel in the partner companies.
- Capture ideas for projects and innovation that can be developed into proposals for feasibility projects.

The other six scoping activities are organized by Strand in WP3, with an additional activity related to the Working Practices Strand in WP1. Each of these activities will identify opportunities for innovation based on the existing background work within each of the Strand areas. These scoping activities are a structured way to allow researchers in the background projects to contribute to the centre's work plan and become part of the centre.

Status of Feasibility Activities The following tables describe the planned feasibility activities for 2016.

SIRIUS PhD Projects at the University of Oxford Several Oxford PhD students are actively engaged in SIRIUS research projects, including Alessandro Ronca, who is working on Incremental Reasoning for Continuous Queries, and Anthony Potter, who is working on Materialisation-based Query Answering. Alessandro's work is still at a relatively early stage (he only started his PhD in October 2015, and since then has had to attend a number of taught courses), but Anthony's work is rapidly maturing, and since returning from a summer internship at Oracle (in their Redwood Shores campus in California) he has produced several important results.

First, he has devised a novel data partitioning algorithm that tries to maximise the "semantic locality" of data by using an approach based on a graph partitioning, where the algorithm tries to create a balanced partitioning while minimising the number of cut edges in the graph.

Second, he has developed a novel query processing technique that exploits the semantic locality of the partition so as to maximise local computation and minimise network communication. The query algorithm runs in a completely distributed manner, with any node able to distribute client queries across the cluster, with no "master" node responsible for accumulating results, and with a novel mechanism for detecting when all answers have been computed. The algorithm also uses a novel approach to message passing and memory management to ensure low and predictable memory usage during query processing, and to avoid non-termination due to full buffers or other memory related issues.

F3: HPC Support for Reservoir Simulation	
Starting date and duration:	TBD (fall 2015 – spring 2016)
Type of Research:	Industrial (feasibility)
Partners involved:	Statoil, Numascale, Dolphin, UiO, Simula
Background/Rationale: To facilitate efficient extraction of oil and gas, and utilize current and future wells in the best possible way, we need extended knowledge about the reservoirs and their behaviour. One way to gain such knowledge is by reservoir modelling and simulation. The accuracy of the simulation results, however, is highly dependent on the granularity and precision of the reservoir model, which again greatly influence the computational power needed to run the simulations. In general, the higher the granularity and the greater the precision of the model, the more computational power is needed. Thus, the most efficient High Performance Computing (HPC) platform is needed to support complex reservoir models, ensure maximum simulation accuracy, and ultimately optimize the reservoir utilization and ensure economically efficient operation.	
Objectives: In this pre-project we aim to identify research challenges to improve the HPC platform support and efficiency for reservoir simulations along three main axes: 1. MPI vs. Multi-threading MPI and Multi-threading (e.g. using OpenMP) represent two fundamentally different parallel programing paradigms. Which of these parallelization techniques that is the most efficient one for a given application domain is not obvious as they both have pros and cons. Along the MPI vs. Multi-threading axis we will explore the two parallel programing paradigms in the case of reservoir simulation, both when it comes to implementation complexity and cost, as well as simulation execution time and efficiency – and identify relevant research challenges. 2. GPU and Accelerator Utilization Due to the tight coupling between processes running in parallel to simulate a reservoir, it is particularly challenging to utilize highly parallel GPUs and accelerators in an efficient manner. On the other hand, GPUs and accelerators (like Xeon Phi) provide excellent watt efficiency and price/performance characteristics. Along the GPU and Accelerator Utilization axis we aim at identifying the key areas of research to help realize the potential of GPUs and Accelerators in the context of reservoir simulations. 3. Cell Convergence in Large Parallel Grid Simulations In large parallel grid-based reservoir simulations, some cells in the grid might converge fast, while other cells require an extended amount of time to finish. These differences in time required for different cells to converge is hard to predict and might lead to poor resource utilization as resources assigned to quickly converged cells remain idle while waiting for other cells to finish. Along the Cell Convergence in Large Parallel Grid Simulation axis we will identify research challenges to improve the software and hardware solutions facilitating large parallel grid simulations to improve overall resource utilization and simulation efficiency.	
Activities current year: Initial discussions along the three axes described above.	

F4: Scalable Data Access for G&G	
Starting date and duration:	2016-01 for 6 months
Type of Research:	Industrial
Partners involved:	Statoil, UiO, Schlumberger, IBM, University of Oxford
Background/Rationale: The Optique EU project has demonstrated the potential for applying Semantic Data Access methods to databases in the Geology and Geophysics (G&G) domain in Statoil. This work is at a prototype stage and has exposed many fundamental and industrial issues that need to be addressed.	
Objectives: Define one or more follow-up activities for the Optique work that address fundamental research problems and industrial implementation issues.	
Activities current year: Review of Optique use case and identification of follow-up fundamental and industrial activities. Preparation of project descriptions, plans and budgets for follow-up activities. Submission of funding proposals for the identified projects.	

F5: Fact Pages for Oil Companies	
Starting date and duration:	2016-03 for 6 months
Type of Research:	Industrial
Partners involved:	Statoil, UiO, Schlumberger, IBM, University of Oxford
Background/Rationale: Ontology-based fact pages provide an effective and maintainable way of allowing personnel to obtain data about assets and equipment. The Norwegian Petroleum Directorate fact pages http://factpages.npd.no/factpages/ are built using this approach and provide a good example of how fact pages work. Introducing fact pages into operational organisations poses both research challenges related to ontology-based data access (OBDA) and practical issues related to data governance, access and structure.	
Objectives: Define and scope an end-user oriented project that will pilot the use of fact pages in one or more oil companies on the Norwegian Continental Shelf.	
Activities current year: Preparation of consortium. Determination of end-user requirements necessary for project definition.	

Finally, all of the above ideas have been implemented in a prototype system based on RDFox, and is being evaluated against several other state of the art systems. At the same time, Anthony is writing a paper describing all of the above work, and we plan to submit this to the International Semantic Web Conference.

SIRIUS PhD Projects at the University of Oslo So far, one PhD student, Vidar Klungre, has started on SIRIUS related projects at UiO. The goal of his project is to analyse and improve the usability of the Optique Visual Query System (OptiqueVQS), and similar systems developed within SIRIUS in the future. OptiqueVQS is an ontology-based visual query system, which allows end-users to easily query for data within their company. It is build as a part of EU-funded Optique Project, but will be extended also after Optique is done in November 2016. More specifically, Vidar will work on the following problems:

- **P1: Adapt interface suggestions based on ontology and data.** The goal of this problem is to improve the suggestions given by the OptiqueVQS even further, by looking at the underlying data. Currently the system only takes the ontology into account, but a good query system should also give suggestions based on the underlying data. The usefulness of exploiting the underlying data in the user interface is inspired by techniques found in standard faceted search systems.
- **P2: Investigate the relation between subconcepts and attributes.** To the end-users, it is not always obvious whether a subconcept is modelled by using ontology suconcepts, or specific attributes. Currently the user interface mirrors the ontology, so the user has to adapt to underlying model, which may contradict with his own model of the domain.
- **P3: Add query support for multiple attribute values.** Currently the OptiqueVQS only supports one filter per attribute, but end-users may want to query for more.

The three problems will be studied in the order they are listed above. Currently, after a half year of work, only P1 has been considered.

6 International Collaboration

Value of Collaboration. International collaboration will contribute to all activities of SIRIUS, including research and development, researcher training, profile building, dissemination, and exploitation. It is easy to see that international collaboration is mutually beneficial:

Value to the Centre	Value to Collaborators
• Raising the profile of SIRIUS • Dissemination and value creation • Access to complementary expertise • Access to a more diverse set of problems and data • Leveraging centre funding	• SIRIUS will be an international leader in research and innovation in a critically important sector of the world economy • Access to the Centre's world leading expertise • Access to laboratory and datasets • Access to industry partners and real problems

Kinds of Collaboration. SIRIUS supports exploratory collaborations with leading groups in "neighbouring" fields, and possibly in other disciplines. The primary mechanism is a *residency program*, funding research visits to/from other internationally leading groups. A particular focus will be on funding longer term visits (of up to 6 months) by SIRIUS's *early-career researchers* to broaden their perspective and contribute to their training and development, and on hosting similar visits. Going beyond such *bilateral academic and industrial collaboration*, SIRIUS also plans to actively engage in the following forms of collaboration:

- *Standardisation.* U. Oxford is a W3C member organisation and played a leading role in the development of the OWL ontology language standard, while UiO have links to the ISO 15926 Oil & Gas ontology effort. These activities provide excellent network building opportunities, and contributes to dissemination and profile building.
- *Research projects* and *Marie Curie networks.* We specifically target Horizon 2020.[15] The combination of research and domain expertise, industry partners and laboratory will make SIRIUS a *very attractive partner* for EU collaborative projects. The Feasibility Project F1, on Next-Generation Maintenance and operations, has resulted in the formation of a consortium, led by SIRIUS, that plans to submit a proposal to the ICT-15 program in April 2016. SIRIUS is also a partner in another application related to cloud computing with the same deadline.

Envisaged Collaborators. SIRIUS' partners have extensive established international networks of collaborators. The prototyping partners are all leading multinationals in the IT area, with global operations. The academic partners have extensive international networks. For example, UiO coordinates the flagship EU funded project Optique, with a consortium that includes several leading European research groups, and as well as participating in Optique, U. Oxford leads a prestigious EPSRC funded "platform" project that brings together academic and industry collaborators including the likes of Google and Stanford.

SIRIUS plans to expand its international network and in particular collaborate with other Centres of Excellence to broaden our research base, gain access to critical expertise, and ensure that our research continues to be internationally leading; e.g., UiO already collaborates with the UPMARC Centre on multicore computing at Uppsala University. It is important to be able to exploit new collaboration opportunities as and when they arise, and as they are indicated by the evolving activities of SIRIUS. Already identified potential collaborations include:

- The *YAGO* project led by Gerhard Weikum at the Max Planck Institute; they are leading experts in the extraction of semantic data from online resources such as Wikipedia.
- The *GLOBUS* project led by Ian Foster at Argonne National Laboratories; they are leading experts in distributed and parallel computing, and are developing tools that allow users to automatically extract/synthesise metadata for diverse collections of data.
- The group of Stefano Ceri at Politecnico di Milano; they are leading experts in stream processing, particularly as applied to the semantic web, and their *ETALIS* system allows ontology-based reasoning to be applied to semantic data in an incremental fashion.
- The *Webdam* project led by Serge Abiteboul at INRIA Saclay–Ile-de-France; they are leading experts on data exchange and distributed query processing.

Alan Turing Institute The activities of the Alan Turing Institute (ATI) continue to ramp up: PhD positions are now being advertised (see https://turing.ac.uk/jobs/2016-doctoral-studentships/), post-doctoral research fellows are currently being interviewed, and faculty research fellows have already been appointed. Amongst the latter is Professor Ian Horrocks, who will act as the liaison and point of contact between SIRIUS and the ATI.

We are currently in the process of organising a workshop at the ATI that will bring together researchers in the areas of semantic technologies, machine learning and statistics, and will provide an opportunity for SIRIUS partners to visit the ATI and meet ATI fellows.

BYTE Project The BYTE project, funded by the European Commission under FP7, runs until 2017, with University of Oslo as one of the partners. The project will, in collaboration with expert stakeholders, develop a vision for big data in 2020 that includes meeting the relevant goals of the Digital Agenda for Europe. BYTE will culminate in the launch of the big data community, a sustainable, cross-disciplinary platform that will implement the roadmap and assist stakeholders in identifying and meeting big data challenges. Furthermore, BYTE will disseminate project findings and recommendations and publicise the big data community to a large population of stakeholders to encourage further innovation and economic competitiveness in Europe's engagement with big data. SIRIUS will devote resources to establish an ongoing relationship with the BYTE project, transfer knowledge between SIRIUS and BYTE and use this to identify scoping and feasibility activities.

7 Recruitment

Mentor Programme The SIRIUS mentor program offers individual SIRIUS researchers a personal development strategy where one person facilitates development of another by exchanging expertise, values, skills, perspectives and attitudes.

The purpose of the mentoring relationship is to gain new perspectives on career development and access to a wider network of resources. It provides the opportunity to build career competence on both sides of the relationship, and gives mutual motivation and inspiration.

Benefits for the mentees are as follows:

- Learn from your mentor's expertise.
- Get feedback in key areas, such as communications, interpersonal relationships, technical abilities, and leadership skills.
- Focus on reflection and actions to further your learning and your career development.
- Learn specific skills and knowledge relevant to professional and personal goals.
- Build relationships and networks for furthering your career.
- Understand organizational cultures and unspoken rules that can be critical for success. Benefits for the mentors:
- Learn from your mentee's background and history to enhance your professional and personal development.
- Gain new insight into the dynamics and learning opportunities of mentoring – also using these skills in your daily work.
- Get satisfaction in sharing your experience and expertise with others.
- Re-energize your career and acquire a fresh perspective on a topic.
- Build a relationship with someone outside your area to expand your network.
- Sharpen your skills for spotting and developing talent.

The SIRIUS Mentoring Programme provides a framework for a mutually beneficial learning relationship. The Programme will contribute to the professional and personal development of both mentees and mentors. It provides training and support, and has the following programme content:

- Each mentor and mentee relationship lasts for 9 months.
- The programme starts in September and ends in June.
- During this period, mentors and mentees are expected to meet 1-2 times a month.
- The SIRIUS mentor program will match mentor/mentee pairs. If a mentoring partnership does not work out, the program management will re-match and reassign mentor/mentee.
- In addition, each year, the programme will organize:
  - Mentor training for mentors.
  - Mentor programme kick-off for new pairs.
  - Two seminars with topics that facilitate personal and career development. These meetings also provide arenas for social interactions between all pairs to share knowledge and expertise.
  - A common closing event.

Mentors are SIRIUS Industry leaders who are motivated to support the career development of a SIRIUS researcher. Mentor are required to fulfil the following criteria:

- Interest in encouraging other people's growth.
- Wish to develop oneself.
- Ability to listen and serve as an inspiring discussion partner.
- Willingness to reserve time for the duties of a mentor.
- It is expected that both mentors and mentees participate in all the activities that are part of the programme.

Recruitment To ensure broad and gender balanced recruitment, we plan educational tracks with both *Master* and *PhD* levels at UiO and NTNU, the two largest universities in Norway, and at U. Oxford, one of the world's leading Computer Science departments. To increase the number of qualified and motivated candidates, a *recruitment group* of promising students at late Bachelor and early Master level will be established, and associated to the Centre through invitations to seminars and events, course recommendations, and opportunities for summer internship and master thesis tuition; by ensuring gender balance in this group, we also aim to encourage female candidates.

Master level: Students are recruited from computer and information sciences at NTNU, UiO, and U. Oxford. Complementing traditional Master studies, SIRIUS also focuses on *executive training*, drawing on the strong platforms for executive education at UiO and NTNU, and *innovation and entrepreneurship*, designed to prepare students for successful careers in innovation related roles in established organisations or new ventures. This is in line with the outward-facing character of SIRIUS and the close interaction between academic and industrial partners.

PhD level: We will recruit from our Master tracks, from industry partners, and through our extensive international academic networks. SIRIUS will cater for both *academic PhDs* (75%) and *industry PhDs* (25%). SIRIUS will establish a PhD school with strong international ties. The PhD school will provide high-quality training and supervision of PhD students through several stages: an *initial stage* which introduces the problem area, defines the PhD goals and an initial progress plan, and establishes a network of supervision and mentoring for the PhD student; an *intermediate stage*, consisting of intensive PhD-level courses with highly qualified lecturers, seminar series, workshops and discussions, regular plenary presentations with feedback on progress and a subsequent revision of the PhD goals and progress plan; and a *final stage* immediately before thesis submission, including a trial defense with realistic feedback.

Number of Positions. SIRIUS has budgeted with fellowships for 25 PhD and 17 Postdoc positions, and 24 person years as Senior Researcher. SIRIUS targets 50% female fellowship holders in the active fellowship positions by the end of the Centre lifetime.

Gender Equality A *Gender Equality plan* will be developed and implemented. SIRIUS aims at an *equal representation* of men and women at all levels through the following core activities.

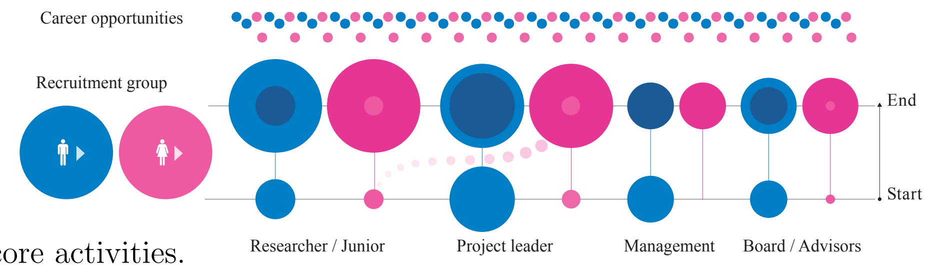

1. **Relevance:** Women tend to seek careers where they see their work making a difference and contributing to better futures. This is why there are relatively few women in ICT — not because of its lack of relevance, but because of its *perceived* lack of relevance. Realising this, we will use the dedicated work package on *outreach* to make visible how the innovations in SIRIUS will have major impact, and how each staff member's effort on SIRIUS contributes to more efficient exploration and safer operations, and thereby a cleaner and more sustainable Oil & Gas sector.
2. **Recruitment:** We aim to achieve gender balanced recruitment in our Masters and PhD programmes by: establishing a *recruitment group* with an equal number of male and female late Bachelor and early Master level students; implementing a radical gender quota system for joining the group; and addressing issues that female students prioritize more than males, including *future job availability*, *relevance of competence*, *international opportunities* and *individual mentoring*. For *external recruitment*, moderate affirmative actions will be applied when hiring.

3. **Career opportunities:** The broad interdisciplinary consortium of SIRIUS results in a *diversity of disciplines and affiliated industries*, and an increased job market for PhD candidates not retained in academia. Furthermore, the HR policies of SIRIUS build on core values like *flexible working conditions*, *job security*, and *family friendly work conditions*, while exploring mechanisms such as *individual carrier planning*, *leadership training*, and *active mentoring*.

8 Communication and dissemination activities

Overview SIRIUS is committed to our dissemination strategy targeting three arenas for dissemination. Towards the *academic community*, SIRIUS promotes aggressive dissemination to invoke and stimulate collaboration with the international research community. The research partners have an outstanding track record of publications in high-quality conference and journals; Prof. Horrocks, e.g., is among the most highly cited scholars in computer science (Google Scholar citation count of over 32,000, as of Feb. 2014). This existing publication track-record provides a springboard for SIRIUS's continued high-quality academic output.

Furthermore, SIRIUS targets *dissemination to industry at large* also beyond the partners in SIRIUS. We will undertake a range of activities designed to ensure the widest possible uptake of our results, and engage with relevant target groups, including industry-based researchers and potential users of SIRIUS technologies. We will present SIRIUS activities and outputs at relevant national and international meetings, industry fairs and industry seminars and arrange tutorials and training courses addressing industry needs. Recent presentations by consortium members include, e.g., the Semantic Technology Conference (SemTech) in San Jose, California; recent industry seminars have been given at, e.g., BAE systems, IBM, Microsoft, Oracle, Samsung and Statoil. We also plan to arrange training courses addressing industrial needs.

Web page, Press and Social Media *Dissemination to the public* is pursued through the popular press (e.g., SIRIUS partners have presented research on several occasions in newspapers and periodicals) and through social media. Twitter will be used for short, timely announcements and as a form of communication with people interested in SIRIUS. A twitter "window pane" will be included in the front page of the website, as well as a SIRIUS blog. The Twitter handle for SIRIUS is @SiriusSfi.

A virtual excellence network will build on on-line communities provided by LinkedIn and the Society of Petroleum Engineers to connect professionals and experts in the field to disseminate SIRIUS results.

BDVA: Big Data Value Association SIRIUS has joined the Big Data Value Association as an academic full member. The BDVA is a Public-Private Partnership (PPP) instrument of the European Commission. Participation in the association's activities will help the centre's dissemination and will also enable the centre to interact closely with European innovation in large-scale data. This membership provides us with a direct channel to influence European policy.

Common Activities with Peripheral Projects Existing peripheral projects, sponsored by the EU and RCN, will be used as dissemination fora and as means to generate project proposals. Current projects are listed in a table on the following page.

Subsea Valley Conference SIRIUS representatives will be holding a master class on scalable data access at the Subsea Valley conference in Oslo in April 2016. This is a first step in what we plan to be a close collaboration between the Centre and the Subsea Valley Cluster. Both the Centre and Subsea Valley share the aim of improving the performance of the Norwegian Oil & Gas sector.

Notes

[1] http://www.idcresearch.com/

[2] http://ec.europa.eu/programmes/horizon2020/

[3] Source: Pål Navestad, ConocoPhillips.

[4] Source: Eldar Bjørge, Statoil Exploration

[5] Zhovtobryukh, Y., Nordkvelde, M. and Reve, T. "Offshore oil and gas as industrial driver", report BI Norwegian business school, 2013

[6] Chesbrough, H. and Apleyard, M.M. "Open innovation and strategy", California Management Review, 31(5), 1998: 32–40

[7] http://www.optique-project.eu/

[8] Baldwin, Carliss, Christoph Hienerth, and Eric Von Hippel. "How user innovations become commercial products: A theoretical investigation and case study." Research policy 35(9), 2006: 1291–1313.

[9] Tichy, Walter F. "Should computer scientists experiment more?" Computer 31(5), 1998: 32–40.

[10] Turner, Lewis, and Nielsen. "Determining usability test sample size." Intl. Encycl. of Ergonomics and Human Factors 3, 2006: 3084–3088.

[11] Walsham, G. "Doing interpretive research". European journal of inf. systems, 15(3), 2006: 320–330.

[12] http://www.gnu.org/licenses/lgpl.html

[13] http://www.apache.org/licenses/LICENSE-2.0

[14] It is nevertheless expected that improved code is voluntarily shared during the lifetime of SIRIUS to improve its total code base.

[15] http://ec.europa.eu/programmes/horizon2020/

Name	Background/Rationale	Objective
Digital oil (Doil)	All phases of upstream oil and gas activities – exploration, drilling, production, abandoning –rely on increasingly information- and knowledge-intensive activities involving information gathering, search, access, analysis and calculations. The work practices of the professional subsurface communities suffer when the IT tools are insufficient, ineffective and/ or inappropriate.	The overall objective of Digital oil (Doil) project is to improve the ability of subsurface communities to harvest the business value from quality data throughout the lifecycles of wells. The five work package objectives are: WP1. Identify existing tactics by production engineers to assess the credibility of sensor-based data in daily operations; WP2. Identify existing user-driven heuristics for generating overviews of the well lifecycle; WP3. Develop a demonstrator with associated work process for situated search of the full lifecycle of well data; WP4. Cultivate a network of subsurface community members, vendors and researchers; WP5. Disseminate results through NTNU Master of Management executive education and international publications in conferences, journals and one edited book.
Optique	Optique, Scalable End-User Access to Big Data, is a large-scale Integrating Project funded by the European Commission under FP7.	Objectives: To design and implement an end-to-end solution to the problem of providing comprehensive and timely access to large scale data sets.
UpScale	Whereas most software today is developed using the object oriented paradigm, multithreaded object-oriented programs do not scale to massively parallel many-core architectures. This project explores alternative mechanisms for parallelization and concurrency control to allow object-oriented programs to scale to this kind of massive parallelism. UpScale is an EU FP7 FET project coordinated by CWI.	The main outcome of the Upscale Project is develop a new programming language that analyses the program to enhance the deployment of software architecture based on ownership types to express the structure, encapsulation and sharing among actors, and session types to express the interaction protocols between various actors. The most important innovation will be the integration of these two technologies, which will be used to express the hand of uniquely references data structures from one party to another along with the exclusive access to shared data structures. In particular, such protocols will involve multiple interaction steps, beyond a simple method call, which justifies the more advanced typing machinery. Extensions to ownership and session type approaches based on lightweight assertions, which can either be checked at run-time or using different verification techniques.
Engineering Virtualized Services (Envisage)	Software is increasingly deployed as services on the Cloud, delivered to meet contracted Service-level agreements (SLA). We need the ability to predict the behaviour and performance of software running on the cloud before it is actually deployed, in order to reason about whether we can meet given service levels and resource management. Envisage is an EU FP7 project coordinated by the Univ. of Oslo.	The overall objective of the Envisage project is to improve the ability of design scalable and resource aware services for deployment on the cloud by means of model-based deployment and analysis. The main outcome of ENVISAGE is a practical open-source framework for model- based development of virtualized services including 1. a behavioural specification language for describing resource aware models; 2. a simulator with visualization facilities; and 3. tool support for automated resource analysis, validation of SLA, code generation, and runtime monitoring of SLA for deployed services.
Scalable Hybrid Variability (HyVar)	HyVar proposes a development framework for continuous and individualized evolution of distributed software applications running on remote devices in heterogeneous environments. The framework will combine variability modeling from software product lines with formal methods and software upgrades, and be integrated in existing software development processes. HyVar is an EU H2020 project coordinated by Santer Reply (Italy). Prof. Einar Broch Johnsen, Univ. of Oslo, is the scientific coordinator of HyVar.	HyVar's objectives are (O1) to develop a Domain Specific Variability Language (DSVL) and tool chain to support software variability for such applications; (O2) to develop a cloud infrastructure that exploits software variability as described in the DSVL to track the software configurations deployed on remote devices and to enable (i) the collection of data from the devices to monitor their behavior; and (ii) secure and efficient customized updates; (O3) to develop a technology for over-the-air updates of distributed applications which enables continuous software evolution after deployment on complex remote devices that incorporate a system of systems; and (O4) to test HyVar's approach as described in the above objectives in an industry-led demonstrator to assess in quantifiable ways its benefits.
Cumulus	This project develops a theoretical foundation for the model-based analysis of resource management and scaling of software running on the Cloud. The Cumulus project is funded by the RCN's FRINATEK programme.	The main goal of CUMULUS is to develop a semantic foundation for static analysis techniques for cloud-aware applications. Cloud-awareness enables the software to negotiate its own quality of service and opens for dynamic and fine-grained resource management. This introduces an element of reflection which goes beyond the state of the art in both formal semantics and static analysis. The project will develop a formal foundation for cloud-aware computing and use this foundation to enable the verification of quantitative assertions about the high-level quality of service and low-level resource requirements of cloud-aware applications.