June 2008

MARKET INTELLIGENCE REPORT

NEXT GENERATION COMPUTING

An initial study of the market for Next Generation Computing systems defined as:
"Systems that enable complex and sophisticated computing applications to be executed in an efficient and cost-effective manner"

CONTENT

EXECUTIVE SUMMARY

This report provides intelligence into the sector defined by ITI Techmedia as Next Generation Computing. In general terms, computing systems enable applications to process information or signals for a purpose, which interconnect with other systems and/or interface with people. In this context, and for the purpose of this report, the definition of Next Generation Computing is:

> ***"Systems that enable complex and sophisticated computing applications to be executed in an efficient and cost-effective manner****"*

This report provides an overview of the Next Generation Computing market, sets out key trends, drivers and inhibitors, reviews the outlook for market development and describes seven market opportunities in greater detail, one or more of which may form the basis for further ITI Techmedia activities in this area.

The objective of the Next Generation Computing foresighting activity exercise was to identify market opportunities which would:

- require more powerful computing systems and/or a larger number of interconnected systems than is being economically realised currently
- include applications which would generate substantially more value from operating more efficiently by being faster, requiring less power or cost
- identify application features not practical or widely implemented today due to a lack of available computing power.

The Next Generation Computing market has been considered both from an application and technology perspective. The scope of the foresighting included applications that would benefit from progress in Next Generation Computing systems, and technologies (both hardware and software) which drive progress in Next Generation Computing and meet the requirements of these applications.

Moore's Law of continuous and ongoing hardware performance gains was mainly driven by generations of faster and more complex monolithic processing units. With power efficiencies rapidly decreasing and physical effects becoming increasingly prominent, the traditional drivers have lost their impetus.

The computer industry continues to innovate to address these challenges. For example, the industry is exploring the use of other semiconductor materials and improved isolators, such as germanium, as a basis for future performance improvements. The industry is also in the process of adopting multi-core architectures to sustain Moore's Law, while semiconductor processes continue to deliver smaller and higher-density structures on silicon.

Whilst hardware continues to pursue performance improvement, software engineering productivity has not exhibited the same rate of performance improvement as hardware. Indeed, with the growing complexity of systems and more and more components working in parallel, overall software development efficiency has decreased.

As with any new markets that offer significant new commercial opportunity, the emergence and growth of aggressive technology start-ups and small companies with innovative products can be expected. Opportunities within the context of Next Generation Computing could represent a significant new and addressable market for technology businesses in Scotland.

Through the creation of underlying technology platforms, ITI Techmedia can play a role in enabling the development of new technology businesses. ITI Techmedia has identified and described seven opportunities that could lead to the creation of technically novel and commercially exciting enabling technology platforms. These opportunities are

- Parallel software tools
- Hardware acceleration
- Natural language processing
- Cognitive radio

- Optical computing
- Systems biology simulation
- Long-tail software services.

Each of these opportunities is described in this Report.

Subsequent analysis has indicated that the following opportunities could form the basis for further ITI Techmedia activity in this area:

- Hardware acceleration
- Optical computing
- Cognitive radio.

Prior to progressing further, ITI Techmedia will seek to engage with ITI Members to validate and confirm those opportunities that should form the basis for its subsequent activities. In addition, ITI Techmedia welcomes R&D programme proposals and expressions of interest in related areas to assist in providing an understanding of the capabilities and commercial interest within Scotland.

1 INTRODUCTION

1.1 Document Purpose

The purpose of this document is to provide a 'snapshot' view of the Next Generation Computing sector in order that the ITI membership:

- has sight of the market analysis activities undertaken in this sector by ITI Techmedia;
- can gain access to market information relevant to the sector;
- is provided with some opportunities that ITI Techmedia will explore further to discover if they offer opportunities for technology innovation that may form the basis of ITI Techmedia research and development programmes in this area.

This document should not be considered as providing a fully comprehensive analysis of the market opportunity space within the Next Generation Computing sector. Such an analysis is beyond the scope of the foresighting activity and this document. This report aims to provide an understanding of Next Generation Computing and its applications, and suggest selected opportunities in this area. It also aims to give those who wish to act as players in this space an appreciation of the market dynamics, scale of the market, innovation potential and timescales.

1.2 Document Structure and Content

This document provides market intelligence for the sector defined by ITI Techmedia as Next Generation Computing (see Section 3 for the definition of Next Generation Computing). The information captured within the document has been obtained following the principles of market intelligence gathering (otherwise known as foresighting) established by ITI Techmedia.

During the process of developing this market intelligence report, both primary and secondary market data were acquired and collated. Primary data was collected by interviewing experts from academia as well as commerce and through a workshop. The primary data gathering process was augmented by desk-based research which was used to obtain secondary data from internationally recognised analysts and other sources. Where possible, the source of the data used in this report has been referenced. This work was carried out with significant contributions from The Technology Partnership (TTP) and the Edinburgh Parallel Computing Centre (EPCC).

The document contains the following sections:

Section 1: Executive Summary. This section summarises the main findings.

Section 2: Introduction. This section covers the background, aims and scope of ITI Scotland. It also provides a high level description of the 'Techmedia' areas of focus. Further background information can be obtained on the website www.ititechmedia.com.

Section 3: Market Overview. This section provides a working definition of the Next Generation Computing market, highlights the main characteristics of the sector and the segmentation of the market. The main trends, drivers and inhibitors are identified.

Section 3: Market Opportunity Assessment. This section provides an analysis of the top market opportunities identified during the foresighting process. For each opportunity it includes a description, unmet market needs, a market overview, the IP potential and timelines.

Section 4: Conclusions and Next Steps. This section provides the key conclusions of the report together with a summary of the next steps that ITI Techmedia intends to take in the area of Next Generation Computing.

1.3 Background: ITI Scotland

1.3.1 Economic Context

A global driver for economic growth is the development and exploitation of technology both for present needs and future requirements. Successful economies are underpinned by a vibrant research base which extends from basic science through to pre-competitive research and development, with a clear focus driven by global market opportunities. Scotland has a reputation for world-class research in many fields and already undertakes significant research activity in several areas which have the potential to provide strong future market opportunities. In addition to the research base, most developed economies have institutes or organisations that promote knowledge generation and increase commercial exploitation capacity. The establishment of such organisations has had significant economic impact over the long term.

1.3.2 ITI Scotland

ITI Scotland is a commercial organisation focused on driving sustainable economic growth in Scotland, through ownership of commercially targeted R&D programmes that deliver world-class intellectual assets.

Specialists from ITI Scotland's three divisions - ITI Techmedia, ITI Energy and ITI Life Sciences - identify technologies required to address future global market opportunities, then fund and manage R&D Programmes and the subsequent commercial exploitation of new intellectual property (IP). This publicly funded company has an active membership programme for interested parties from the business, research, academic and public sectors. Members enjoy exclusive access to market foresighting (such as that contained within this report), the opportunity to participate in leading-edge R&D Programmes and networking opportunities brought about by regular meetings of a growing network of like-minded organisations.

The divisions also interact with each other to identify potential overlap or "white space" market opportunities between ITI Techmedia, ITI Life Sciences and ITI Energy.

They are a centre or hub for:

- identifying, commissioning and diffusing pre-competitive research that is driven by an analysis of emerging markets;
- managing intellectual assets to maximise commercial and economic value for Scotland.

An active membership is core to the ITI Scotland model. Membership is open to companies and research institutions willing to participate actively in its activities. ITI strategy and operations are actively guided and supported by Members. ITI Scotland seeks Members with a broad global perspective on markets and new technology directions, as well as a local focus, to ensure that propositions will be transferred effectively into the Scottish economy.

1.4 Definition of the Techmedia Sector

ITI Techmedia is centred on the development and creation of commercial opportunities encompassing information and communications technologies and digital media sectors. The activities of ITI will bring Scotland's economy to the cutting edge of emerging markets by allowing local companies to access and build upon pre-competitive technology platforms developed by ITI.

The term 'Techmedia' arose out of the need to reflect the market evolution of information and communications technologies and digital media. The overall trend in the marketplace is governed by a value chain ranging from content/application generation through delivery to consumption. Content, service provision, delivery channels and enabling and managing technologies can no longer be treated in isolation and ITI Techmedia seeks to operate throughout the value chain.

The key to identifying opportunities for research and development lies in a process called "market foresighting" which involves detailed market and technology analysis to identify trends, evolving requirements and potential demand for new technology. ITI Techmedia compiles the output of this activity into market intelligence reports which are published to Members. This market foresighting informs our R&D Programme identification process.

The Techmedia sector is potentially very broad. Hence a phased approach to market foresighting has been adopted. Previous foresighting has concentrated upon a number of major market areas:

- (Remote) Health
- Commerce and Finance
- Learning and Education
- Communication Services
- Entertainment and Leisure
- Digital Cinema
- Nanotechnology
- Homeland Security
- Ubiquitous Computing

To date, these foresighting activities have helped ITI Techmedia to identify a number of R&D Programmes:

- ***Games-Based Learning*** - to develop a differentiated creation and authoring platform to simplify the creation of games-based learning content. Completed January 2007.

- ***Machine-Readable Security Tagging*** - to develop an end-to-end system solution, featuring a range of component technologies required to protect brands and combat the growing global threat to products from illegal counterfeit activity. Completed March 2007.

- ***Ultra-wideband Wireless Communications*** - to develop the components, system and network management elements for ultra-wideband wireless technology in consumer markets.

- ***Condition-based Monitoring*** – to apply sensors and networks technology to condition-based monitoring for predictive intervention in animal health matters.

- ***Biosensors*** - to create a technology platform that will facilitate both diagnosis and treatment of infectious diseases.

- ***Online Game Development*** - to establish a world-class development platform to facilitate the efficient production and distribution of online PC, console and hand-held games.

- ***Backlighting Using Polymer Optics*** - to develop a novel backlight platform for liquid-crystal flat-panel displays to improve viewing quality, reduce weight and improve power efficiency at lower cost.

- ***Software Integrity Engineering*** - to develop novel, user-friendly, code design and development tools that assist in the identification and elimination of critical, high impact, software errors for mainstream applications.

2 MARKET OVERVIEW

Over the next 10-15 years novel computing platforms that offer optimal price/performance in support of a plethora of new applications are expected to emerge for use in both the home and the workplace. This market is technologically complex, and these new platforms will be enabled by significant enhancements in the underlying technologies used.

This section presents an overview of the market definition and segmentation for Next Generation Computing and outlines a high level market structure and major market drivers and inhibitors that could influence the market development.

2.1 Market Definition

In general terms, computing systems enable applications to process information or signals for a purpose, which interconnect with other systems and/or interface with people. For the purpose of this report, the Next Generation Computing market is defined as follows:

> ***"Systems that enable complex and sophisticated computing applications to be executed in an efficient and cost-effective manner"***

The scope of foresighting includes:

- applications that would benefit from the availability of Next Generation Computing solutions
- technologies, both hardware and software, that enable the creation of Next Generation Computing solutions that meet the necessary performance requirements of applications. These performance requirements may include faster calculation speeds, power reduction or reduced cost versus currently available technologies.

Specific examples in the area of Next Generation Computing include:

- High-performance computing (HPC) (e.g. massively parallel and distributed) for the execution of applications involving very complex algorithms or the processing of huge amounts of information in a timely fashion
- Embedded computing that enables the support for sophisticated applications that require powerful processing capabilities at preferentially low power in a device
- Appropriate software tools, suitable for new and evolving computing architectures such as tools for the development of applications that require robust and reliable concurrent code execution.

2.2 Market Segmentation

The area of Next Generation Computing is subject to both 'market-pull' and 'technology-push'. As a result, the market for innovative and valuable solutions has been segmented from both an application and an underlying technology perspective. Technologies are considered from both hardware and software aspects.

2.2.1 Segmentation by application

The Next Generation Computing application space is extremely broad. In order to provide some indication of the range of applications that may be affected by Next Generation Computing technologies, the application areas are grouped into the following broad areas:

- Science and engineering, encompassing activities ranging from novel science through to manufacture
- Business/enterprise, encompassing various aspects of the supply chain from market proposition development through to end user consumption
- Personal devices and computer platforms, including network-centric capabilities
- Military, encompassing both field use and training.

For each of these application areas, prominent examples of applications that may benefit from Next Generation Computing technologies are listed below.

Science and engineering:

- Analysis of experimental data
- Exploratory research
- Biological systems
- Biotechnology
- Pharmaceutical research and development
- Chemistry
- Medical
- Physics
- Geophysical engineering
- Energy
- Semiconductor IC design
- Computational fluid dynamics (CFD)
- Finite element analysis (FEA)
- Automotive
- Space
- Environmental

Business/enterprise:

- Tracking, logistics
- Next generation internet
- Games (server-side)
- Digital rights management (DRM)
- Professional imaging (medical imaging, 3d modelling)
- Content adaptation (mobile, online)
- Robotics
- Driver assistance, driverless vehicles
- Transportation infrastructure
- Security
- Industrial systems
- Sensor networks, e.g. for monitoring environmental or industrial facilities
- Mobile infrastructure
- Market analysis
- Customer relations management (CRM)
- Enterprise resource planning (ERP)
- Virus protection
- Mobile communications

Personal devices and computing platforms:

- TV including IPTV, High Definition, and 3D
- Gaming (client-side)
- Multimedia
- Graphics
- Mobile devices
- e-Butler, intelligent agents
- Biometrics
- Human-computer interface (HCI)
- Artificial intelligence (AI)

Military:

- Conflict scenario simulation
- Training simulation
- Weapon simulation

2.2.2 Segmentation by hardware technology

The range of technologies that can be utilised in the creation of Next Generation Computing hardware solutions is extremely varied. In order to provide some indication of the range of applicable hardware technologies, they are grouped into the following broad areas:

- Systems technology that provides the framework for the Next Generation Computing technology system, encompassing technologies ranging from devices to global systems
- Processor technology, encompassing various technologies and techniques for the execution of machine instructions
- Embedded and mobile technologies, encompassing those technologies that will enable the novel embedded and mobile computing applications at optimal performance.

For each of these areas, prominent examples of technologies that may enable the creation of novel Next Generation Computing solutions are listed below.

Systems technology:

- Memory/bus architectures
- Massive server farms
- Distributed computing architectures, e.g. grid, cluster, peer-to-peer
- Adaptable computing, wide area resource scheduling
- Flexible connectivity
- Novel HPC architectures

Processor technology:

- Multi-core processors
- Hyper-threading
- Many core
- Mixed core
- Reconfigurable hardware, e.g. FPGA[1]
- XML[2] hardware
- GPU[3]
- Other hardware accelerator technologies
- Memory chips
- System on chip (SoC, processor + memory + communication)
- Interconnection of core processors
- Resource sizing (e.g. cache)
- Single-processor architectures, e.g. SIMD[4], MIMD[5], Vector
- Power efficiency
- Signal processing
- Optical computing
- Quantum computing

Embedded and mobile:

- Mobile computing
- Broadcast
- DSP[6]-based systems
- Reconfigurable hardware
- Intelligent networks
- Shared processing
- Multi-core processors
- Microcontroller systems

2.2.3 Segmentation by software technology

A broad range of technologies can be utilised in the creation of Next Generation Computing software solutions. In order to provide some indication of software technologies, they are grouped into the following broad areas:

- Software architecture technologies that enable the creation of an abstract representation of a software system
- Software development technologies that enable the efficient and error-free creation of code

For each of these areas, prominent examples of technologies that may enable the creation of novel Next Generation Computing solutions are listed below.

Software architecture technologies:

- System-level design
- Service-oriented architectures (SOA)
- Virtualisation

Software development technologies

- XML processing
- IDE[7] (parallelisation) support

Also, a number of **specific software technology components** have been identified, such as:

- Artificial intelligence including natural language processing
- Simulation technologies

[1] Field Programmable Gate Arrays

[2] Extensible Markup Language

[3] Graphics Processing Unit

[4] Single Instruction, Multiple Data

[5] Multiple Instruction, Multiple Data

[6] Digital Signal Processor

[7] Integrated Design Environment

2.3 Characterisation of the Ecosystem

Although some next generation computing systems share certain common technology components and distribution models, when looking at such broad market segments there are significant differences in ecosystems depending on the application type and area. For example, software applications for mobile and super-computing platforms have very different routes to market. It is therefore challenging and difficult to describe a generic ecosystem that suits all the various segments.

To facilitate understanding, a high-level value chain is described in Figure 1 below and illustrated by an example of the mobile telecom industry (consumer segment). More specific descriptions of individual market structure and relevant key players are provided in the following market opportunity section.

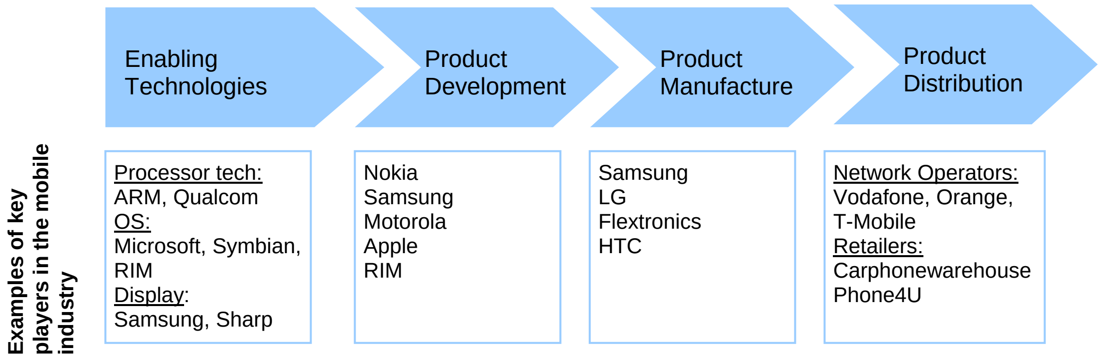

Figure 1: High-level value chain of the mobile industry [Source: ITI Techmedia]

There are four main dimensions presented in this high-level value chain.

- **Enabling technologies** from key technology providers enable key functions and services for end users. Such key technologies are usually well protected and often licensed to component suppliers or product manufacturers. Enabling technologies could encompass a broad variety. Looking at the mobile handset industry for example, enabling technologies include core processors (major providers include ARM and Qualcomm), operating systems (major providers include Microsoft, Symbian, Google and RIM), as well as mobile displays (major suppliers include Samsung and Sharp).
- **Product development** comprises companies in the value chain who define and develop products. They are mainly focusing on the management of the product lifecycle from market requirements capture, through product development, to commercialisation. Product components or technologies are either sourced from third party suppliers or developed in-house. In the mobile industry, examples of players include Nokia, Samsung, Motorola, Apple, and RIM etc. Close partnerships with enabling technology providers are often required to develop mainstream products.
- **Product manufacture** comprises companies who make products based on product designs. Product developers can also be product manufacturers if they have in-house manufacturing capabilities, or outsource manufacturing to third parties. A good example is the fabless semiconductor industry where companies like ARM and CSR outsource the entire manufacturing business to third parties like TSMC and Flextronics, allowing them to focus on technology development and product design. Product manufacturers often work closely with key technology providers for key components and product designers/developers for finished products. In the mobile industry, examples of product manufacturers include Samsung, LG, Flextronics and HTC.
- **Product distribution** uses distribution channels to sell products to end users. Distribution channels are crucial to the success of a product. Companies invest significant resources in marketing, promotion and partnership in order to successfully market their products. Some product developers such as Dell used a direct sales model selling products directly to customers without any intermediaries, while the majority of businesses use distributors and/or retailers to reach potential customers. In the mobile industry, the major distribution channels for mobile phone manufacturers are the network operators e.g. Vodafone and Orange. Network operators have a very strong

position in the value chain which allows them to request phone makers to build phones following their specifications. Retailers like Carphone Warehouse and Phones4U also play an important role in distribution.

For some segments in a value chain, the market is dominated by a handful of big players. For example, the desktop/server CPU market is dominated by Intel and AMD, and the desktop OS market is dominated by Microsoft. Big dominant players have huge influence on the overall market development. Entry barriers are high for new market entrants requiring significant upfront investment and/or disruptive technologies. In other segments, the market is fragmented where many smaller players are active and new entrants often emerge. Examples of such markets include data analysis and power management.

2.4 Market Trends, Drivers and Inhibitors

Key trends, drivers and inhibitors impacting the area of Next Generation Computing are identified below.

Key Trends

- Moore's Law in semiconductor processes – the gate density is steadily growing, and processors costs are declining, enabling new applications with higher processing power
- Continued progress in hardware performance - mainly through multi-core, parallel processing
- Rising connectivity and number of connected devices – in general purpose computing as well as appliances and sensors/actors

Key Drivers

- Significant technology progress in both computing and communications
- Multi-core processors are becoming available and affordable
- Requirement for reduction in energy consumption for all computing platforms with a view to reducing carbon emissions and operating costs
- Increasing performance demand from increasingly complex applications, e.g. in multimedia processing

Key Inhibitors

- Software engineering productivity not progressing in line with growing complexity of interconnected systems and applications
- Standardisation and regulatory barriers, particularly related to the use of radio spectrum
- Large legacy systems slow down the migration to novel systems

2.5 Next Generation Computing Market Outlook

In the more general sense next generation computing comprises a broad range of systems that enable efficient execution of increasingly complex and sophisticated applications. It covers a great breadth of vertical markets, such as medical, energy, online game etc. The overall market opportunity related to Next Generation Computing is huge.

Globally, the market for information and communications technology (ICT) was worth some EUR2.23 trillion in 2007[8]. Europe accounts for around one third of the global ICT market opportunity, and the value of the European market for:

- software and IT services was expected to grow from EUR207 billion to EUR231 billion between 2007 and 2009
- IT equipment was expected to grow from EUR93 billion to EUR95 billion between 2007 and 2009

[8] Source: EITO in collaboration with PAC, Idate, NetConsulting

- telecommunications equipment was expected to grow from EUR71 billion to EUR74 billion between 2007 and 2009
- digital consumer electronics equipment was expected to remain at around EUR61 billion in 2007 and 2009

At the technology segment level, the size of individual segments can also be huge. Taking the multicore processor market for example, the shipments of multi-core microprocessors are expected to grow from 14.8 million in 2005 to 638 million units in 2015[9]. Market revenue is expected to increase from USD2.6 billion to USD64.8 billion accordingly.

It remains a significant challenge to provide a meaningful market forecast for such a general market as Next Generation Computing. In addition to the ICT market data above, the sections relating to the seven specific market opportunities identified present more specific market data.

[9] Source: iSuppli

3 MARKET OPPORTUNITY ASSESSMENT

The objective of the foresighting process is to identify areas of opportunity for the development of technology platforms which address unmet market needs that:

- are unlikely to be satisfactorily addressed by current solutions or approaches in the short term
- have the potential to be addressed in the medium term by new, technology-based developments
- are likely to provide a significant revenue opportunity for market participants in the medium term.

This section describes seven specific market opportunities for Next Generation Computing and the methodology used to identify these opportunities.

3.1 Methodology

Potential market opportunities for the application of Next Generation Computing solutions were identified through research and analysis. The objective of this exercise was to identify market opportunities that would:

- require more powerful computing systems, with a larger number of interconnected systems, than those available currently
- include applications which would generate significant incremental value from running faster and cheaper, serve more clients/customers and achieve higher throughput than is possible currently
- use applications features which are not practical or widely implemented today due to a lack of available computing power.

As a result, some 45 potential market opportunities were generated, as listed below:

1. Java Virtual Machines (JVMs) for parallel architectures
2. Open Multi-Processing (OpenMP) programming tools for Java
3. Multi-core software development tools
4. Programming for parallel computing
5. Parallel language compilers and tools
6. Automatic memory pattern analysis
7. Data/memory access in parallel computing systems
8. FPGA compiler tools for high-level languages
9. Dynamically adaptable hardware platforms
10. General Purpose GPU (GPGPU) compilers and tools
11. Lightweight operating systems for multicore
12. Optimisation of C++ template code
13. Grid computing
14. Image, audio or video search by example
15. Media search engine
16. HCI - analysis of brain signals - brain interfaces
17. Novel recommender engines
18. Personal/professional video creation, editing, mash-ups
19. Systems biology simulation
20. HCI - mobile listen and cues
21. Mobile - software/cognitive radio
22. Software-defined radio (SDR)
23. Optical computing and communications
24. Optical encryption
25. Automated medical image analysis

26. HCI - high-resolution wall-size displays
27. Handling large datasets
28. Software services
29. HCI - mobile human senses augmentation, e.g. augmented vision
30. Realistic real-time graphics generation
31. Natural language processing
32. Speech recognition (talker independent)
33. 3D reconstruction
34. Engineering and maintenance of large heterogeneous systems
35. Self-repair/self-healing computing infrastructure
36. Power management solutions for many core parallel computing systems
37. Power saving by employing parallel computing (embedded)
38. Virtual ID management platform
39. Real-time information transfer between real and virtual world
40. System Level Design – to determine hardware/software partitioning
41. Massively distributed systems
42. Biological computing
43. Data visualisation tools for multi-factor analysis
44. Unstructured text search in enterprise applications
45. Mobile wireless communications - pervasive wireless infrastructure

Opportunities in the long list have been evaluated and prioritised, and a subset of opportunities have been selected for more detailed investigation and validation. The opportunities were assessed against a set of criteria, to identify the most promising opportunities. The assessment criteria used were:

- Commercial potential
  - size of addressable market and potential market share
  - existing competition and potential threat through new entrants
  - market window
- Innovation potential
  - ability to meet unmet needs
  - technology maturity level
  - potential for protecting IP
  - size of investment required
  - time frame for development
- Next Generation Computing potential/applicability i.e. would Next Generation Computing enable this opportunity?

As a result, seven opportunities with the highest ranking were selected for further analysis, and were further validated through discussions with experts (e.g. telephone interviews, workshop). The selected opportunities are listed below together with a brief description of the market proposition.

The rest of this section describes in detail the top seven selected opportunities as well as giving brief details of the next six opportunities that were favourably assessed but not selected for further analysis.

3.2 Description of Top Market Opportunities

The work described above highlighted several areas of potential opportunity with significant scope for innovation and addressing unmet needs. These areas represent an opportunity for the creation of valuable IP for markets developing within a 3-10 year timescale. The short-listed opportunities are summarised in Figure 2 below.

Opportunity	Summary
Parallel software tools	Integrated development environment and/or runtime tools to help create more efficient computing systems more effectively, taking advantage of multi-core and distributed infrastructure, and exploiting data and code parallelism
Hardware acceleration	Tools that design/(re-)configure hardware accelerators to optimize runtime performance of specific applications; tools that partition the code of applications in such a way, that specific applications run optimally on a combination of standard processors and hardware accelerators
Natural language processing	Algorithms/ tools that enable machines to understand natural language
Cognitive radio	Utilize better the available spectrum through flexible/agile radio sub-systems and intelligent signal processing and decision-making (enabled by sophisticated software), to improve bandwidth, resilience to interference, availability of services
Optical computing	Development of optical transistors, logic elements and further integrated circuitry, to create more powerful computing nodes than can be achieved today with electronics technology; apply optical computing technology in the first place to realise completely optical communication systems, including board-level interconnect, as well as optical routers and switches
Systems biology simulation	Developing simulations predicting the efficacy and side-effects of candidate drugs; upside in predicting outcomes of other medical interventions
Long-tail software services	Tools and techniques to customise software and data in a way that customised solutions (the "long-tail" of software) have a better fit to user requirements and are more valuable for users. Customisation encompasses data (e.g. user-generated content "mash-ups", "enterprise mash-ups") and code (e.g. componentisation, semantic description). Making customised solutions available as software services enables much better protection of the solutions against illegal copying, and enables more favourable business models

Figure 2: Short-listed market opportunities [Source: ITI Techmedia]

Each opportunity is described further according to the following structure:

- opportunity description
- identified un-met needs
- market overview
- IP potential
- timescales (market window)
- conclusions.

3.3 Parallel Software Tools

3.3.1 Opportunity Description

Over the past three years, the move to an increasing number of processing cores has been a clear trend in microprocessor design, spanning high-performance computing, desktop, mobile phones and communications. Yet the development approach preferred by the vast majority of software engineers is still towards linear program flow, making poor use of multiple cores.

In the most general terms, this opportunity is about providing tools that would achieve one or both of the following goals:

- allow legacy, sequential application software, libraries and/or template code, to be parallelised with a high degree of automation[10] and minimal modification
- make it easier for today's programmer community to make more efficient use of multi-core and distributed computing systems.

The focus of this opportunity is on tools to develop applications more efficiently on multi-core, general-purpose processors e.g. being developed by Intel, AMD and ARM; more specific and reconfigurable hardware, such as FPGAs, was considered as a separate opportunity (see the section on hardware acceleration). Tools should be accessible to software programmers of standard ability, not restricted to the much smaller number of hardware/software engineers who are comfortable working with parallel hardware platforms with existing, low-level tools.

For the first of these goals, a key requirement is that the programming languages must not change: it is considered too expensive if the software has to be partially rewritten, or manually analysed line-by-line, for example to add compiler hints. This is a "holy grail" that many researchers and corporations have struggled to achieve over the last two decades, with billions of dollars of investment. The main successes have focussed on instruction level parallelism, the use of compiler directives and the addition of simple loop-based language constructs to existing languages.

For the second goal, the range of possible approaches is much wider, potentially encompassing new languages, compilers, parallel libraries, or other tools such as debuggers and profilers. Some prospective areas for R&D were not pursued further as they would be difficult for ITI Techmedia to address, either because of limited IPR potential or the presence of strong existing players already targeting the market, for example:

- development of a new parallel language, or language extensions to existing programming languages (similar scenario to Sun's commercialisation of the Java language or the US government's DARPA-funded HPC initiative to develop new supercomputing languages);
- development of new compilers to part-parallelise or assist with the parallelisation of applications such as those available from companies like Rapidmind and Codeplay;
- development of runtime libraries to parallelise common functions.

One of the strongest propositions is thought to be the extension to multiple cores of the "write once, run anywhere" concept of Java and (more recently) Microsoft's C# and .net tools. These compile software into an intermediate language known as a bytecode, which can then be executed without modification on a wide variety of hardware. This is achieved by running the bytecode through a hardware-specific bytecode interpreter or runtime system, which translates the bytecode instructions into the system's machine code, and interprets calls to the operating system and other standard functions as required. Through the use of techniques such as just-in-time (JIT) compilation and Java-accelerated CPU instruction sets, bytecode languages can approach or in some cases exceed the performance of traditional compiled languages such as C and C++. In its simplest incarnation this proposition could be to develop a standard-compliant Java runtime/JIT compiler that would be able to take advantage of multiple cores.

Other opportunities may arise with model-driven software approaches, where tools could support core-level parallelisation of software generated through a modelling language such as Unified Modelling Language (UML), or targeting higher-level languages such as that used in Matlab.

---

[10] full automation is a long-term goal

Propositions

(1) To provide development and runtime tools (compilers, profilers, runtimes, libraries) to help create more efficient computing systems in a more automatic way, taking advantage of multi-core and distributed infrastructure, data and code parallelism.
(2) To provide tools around a new intermediate language (for example a bytecode):
   a) Tools and algorithms supporting parallelisation of the bytecode, and that can be compiled from software written in existing popular languages such as C++, C# or Java. Interestingly, this approach may also fit well with the need for better FPGA-programming tools discussed in the next section.
   b) To develop parallel runtime interpreters/JIT compilers for bytecode languages, in particular Java and C#, that would be able to adapt automatically to the available hardware and make use of multiple cores.

Applications

The results of R&D efforts in this area would be expected to lead to the creation or incremental development of products such as:

- Software analysis and development tools for existing languages e.g. C/C++, C#, Java.
- Standard thread, runtime, template libraries for software development.
- Parallel runtime interpreters/JIT compilers for byte-code languages.

In general, parallel software tools support the efficient creation of computationally demanding code, which at present can only be achieved on the basis of many-core processing systems. Examples of demanding applications are e.g. simulations and real-time video processing.

3.3.2 Un-met Market Needs

The following un-met needs and requirements have been identified:

- enabling significant numbers of software engineers able to "think in parallel"
- enabling typical software engineers to produce efficient parallel code
- existing development challenges for single processor code are multiplied by the need to consider parallel processing target platforms
- ability to capture and debug dynamic faults
- efficient maintainability of parallel code
- integrated development environments including parallel tools

3.3.3 Market Overview

Market structure and key players

The following figure provides an example of the structure of the software development tools market, targeting the developers and integrators who are the main customers of these tools. Examples of the product offerings and commercial players in the market are also given.

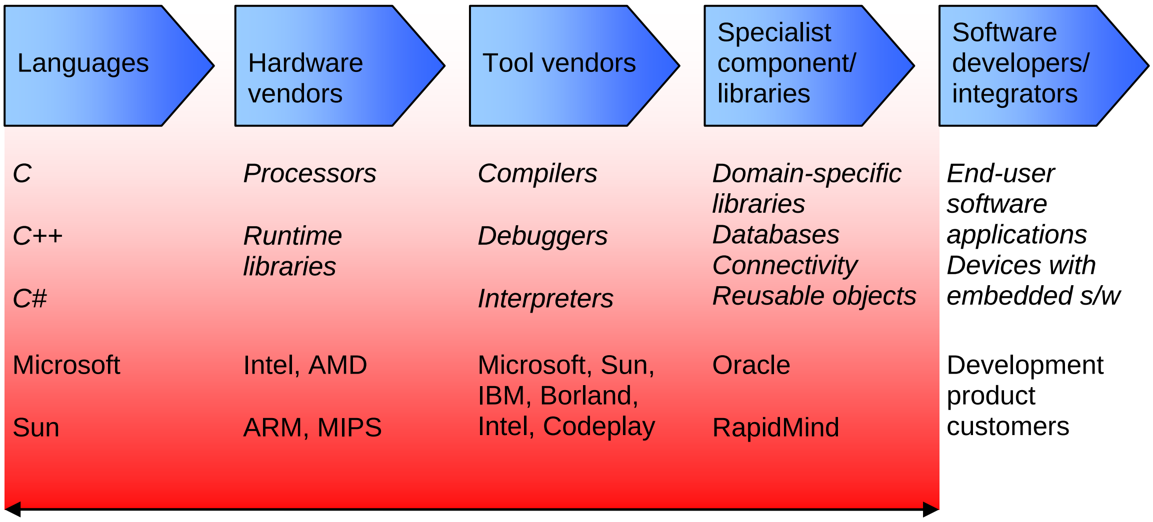

Figure 3: The structure of the software development tools market [Source: ITI Techmedia]

Market size

Software development tool licensing is a multi-billion dollar global market. For example, there are over 8 million developers using Microsoft's Visual Studio tools, which target professionals and cost from about USD500 per licence.[11] Other popular tools include free/open source and commercial variants, from companies/organisations such as Borland, IBM, Sun and the Eclipse Foundation. However, although this market segment is large, players other than Microsoft have struggled to make significant money from it. In particular a trend has emerged for tools such as Eclipse (into which IBM injected much of its compiler and development tool IPR) and Sun's Netbeans to be made available for free, leading to a decline in total revenues.[12]

Intellectual property value in the languages market segment has also proved difficult to realise. Novel programming languages similar in nature to C and C++ are very difficult to patent, and mainly offer such value through training/education programmes, and less so via compilers and development tools. Two more recent languages have become generally successful, though through much more integrated offerings: Sun's Java language is now widely used in applications from enterprise software to mobile phone user interfaces, while Microsoft's C# language – created in response to Java – is also very popular, especially for server application development. Sun makes money from Java through licensing development tools, reference implementations and a portfolio of patents in the runtime, while Microsoft uses C# to sell development tools and operating systems. Neither Sun nor Microsoft gives figures, though the estimate of Java revenue to Sun is in the order of USD10 million per annum (2006).[13]

In contrast, the market for high-performance computing (HPC) systems including supercomputers, specialist components and libraries is more buoyant. Systems serve a broad range of applications, with a wide range in unit costs (extending from tens to millions of dollars). HPC systems are manufactured in lower volumes, and software frequently relies on free components such as the LAMP (Linux, Apache, MySQL, Perl/PHP/Python) software stack. Nevertheless, there is significant value across the market, with IDC estimating that the technical computing market (hardware and software) exceeded USD10 billion in 2006, with revenues 9% greater than the previous year.[14] Other attractive target markets could be computer games, mobile applications and business transactions.

The route to market is expected to be through one or more of the following companies.

---

[11] http://www.regdeveloper.co.uk/2007/06/09/vs_shell_eclipse/ (June 2007);
http://www.microsoft.com (Jan 2008)

[12] http://www.idc.com/getdoc.jsp?containerId=207441 (June 2007)

[13] http://www.forbes.com/global/2006/0522/072.html.

[14] http://www.idc.com/getdoc.jsp?containerId=209251 (Nov 2007).

- Software vendors (e.g. Microsoft, Sun, Metrowerks) with existing software development tool-chain
- Independent specialist vendors, for example as a plug-in to the Eclipse framework[15]
- Major developers/resellers of multi-purpose software or software-enabled services (e.g. IBM, Google, Nokia)
- Processor vendors migrating towards dense multi-core (e.g. Intel, AMD, MIPS, ARM, Marvell, NEC).

Enablers

- Parallel development tools have the potential to add value to existing offerings, allowing tools and developers to keep pace with hardware developments. This is particularly the case with tools targeting general-purpose hardware such as PCs, though also may apply to large specialist markets such as mobile phone platform and application development.

Barriers

- Technically, it has proved very difficult to take advantage of multiple cores – in particular a variable number – using normal sequential software: as one of our interviewees commented "there is no such thing as a parallel compiler".
- Two alternatives provide a potential threat to the value of the proposed tools:
  - software using multiple threads of execution, a technique that is long-established and becoming more widely used by developers, and that scales better to multiple cores;
  - following the Pareto rule to optimise manually those sections of code that can best take advantage of parallel hardware. The latter technique is often favoured by vendors of specialist hardware and was cited by several interviewees.

3.3.4 IP Potential

There are mainly two routes for protecting rights in software intellectual property – patents and copyright:

- New bytecode language, and/or runtime interpreters/compilers for Java/C# – patents in the methods to create and decode the bytecode, and its implementation on multi-core hardware; licensing revenue from the bytecode specification and reference implementations. In the longer term, patents and know-how on optimising multi-core hardware to make better use of the language(s).
- Development and runtime tools (compilers, profilers, runtimes, and libraries) – niche market(s) would most probably have to be targeted, with potential to create copyright in the developed products and possibly application-specific patents.

Generally the patenting system has not provided for software protection, being both very slow and cumbersome and many regimes not supporting it at all. Software companies simply rely on trade secrets and booby-trapped code, for implied copyright, to protect their vested interests and hope parallel discovery will not happen during the short market window associated with innovation in this technology sector.

Recent UK High Court cases, in particular Aerotel Macrossan (November, 2006) and Astron Clinica (March, 2008) have sought to clarify the boundaries of patentability of software and how the exclusions in the patents act will be interpreted. The former resulted in confusion and a departure from current European practices, with the following Astron Clinica retrial seeking to redress the balance.

---

[15] http://www.eclipseplugincentral.com

3.3.5 Timescales

The market for these propositions already exists: quad-core processors are now available from vendors such as AMD, ARM and Intel, dual-core processors have already become widespread in desktop PCs, while high-performance computing workstations and supercomputers offer 8 or more cores (sometimes even thousands) today. It is expected that, by 2011, a majority of new desktop PCs will have two or more cores, while high-end machines may have 8-16.

Markets for niche (domain-specific) runtime libraries for parallel architectures are now starting to be addressed, by CPU vendors such as Intel and a small number of independent software houses although many market opportunities still exist. Similarly, the broader market for development and runtime tools is served by leading players who are obviously aware of the trend towards multi-core. New companies, such as Rapidmind, are springing up to take advantage of the current opportunities. Thus it is likely that this market will become more difficult to enter about 3 years from now, however no step-change solutions are visible today.

The market for runtime interpreters/compilers for existing bytecode languages (Java, C#) is expected to be on the same one to three year timescale, as Microsoft and Sun need to address the multi-core development challenges before potential competitors or new entrants do. Google, with its Dalvik alternative to conventional Java virtual machines and its deep in-house knowledge of distributed computing, would be particularly well placed to disruptively enter this market.

Bringing a new language to market, despite the dire need motivating such an approach, will be much slower due to the need to train up programmers as well as the massive task of creating a fully-integrated reliable platform. The timescales of introduction of C# and Java both indicate that this could be expected to take five to ten years and require a significant upfront investment.

3.3.6 Conclusions

With the increasing number of processing cores, there is a clear need to provide tools support for application developers to produce efficient parallel codes. Monetisation of parallel software tools would require careful consideration.

3.4 Hardware Acceleration

3.4.1 Opportunity Description

In the context of this report hardware acceleration is defined as (elements of) hardware designed to act in support of a traditional microprocessor to accelerate computation.

For the past 30 years the microprocessor has driven advances in computing technology – delivering higher and higher levels of computational performance from ever faster and more complex silicon designs. Moore's Law, which states that the number of transistors that can be inexpensively placed on an integrated circuit is increasing exponentially, doubling approximately every two years was first coined in 1965 by Gordon E. Moore of Intel. Although many commentators have spoken recently of the end of Moore's Law, this is not the case – the exponential increase in processor clock-speeds and processor complexity may have ceased but the number of transistors on a single silicon chip has continued to increase as before through the advent of multi-core architectures.

As discussed elsewhere in this report, the growth of multi-core architectures cannot lead to a linear increase in computational performance due to memory bandwidth and programmability issues (see also Amdahl's law[16]). Therefore, in the past five years, a number of new hardware acceleration technologies have been brought to the market, each of which claims to deliver large computational performance gains from relatively inexpensive hardware generally packaged as a plug-in/add-on to a microprocessor-based motherboard.

---

[16] http://en.wikipedia.org/wiki/Amdahl's_law

To date, no single technology has delivered a full solution across the general numerical computing sector, and thus opportunities for innovation exist in a variety of areas.

There are a number of different acceleration technologies on the market today ranging from highly flexible Field Programmable Gate Array-based devices (FPGAs) through General Purpose Graphics Processing Units (GPGPUs) to Digital Signal Processors (DSPs) and highly bespoke accelerators.

Compared to traditional microprocessors, they all share the characteristics of low power and low real estate requirements per operation. This makes hardware accelerated computing solutions an appealing opportunity for those with power constraints (e.g. modern data centres in the banking sector) or space constraints (e.g. on board aircraft).

FPGAs

Field Programmable Gate Arrays (FPGAs) have been used in embedded computing for many years but only recently become serious contenders in the general computing space. The first commercially available FPGA-based reconfigurable computer was developed in Scotland in the late 1980s[17] and Scotland is seen as a global centre for the technology.

An FPGA consists of a large number of logic devices which can be configured together in any way to process the data which is presented to them. The embedded logic devices range from simple logical gates to DSPs, floating point units and full microprocessors. The keyword associated with FPGAs is flexibility – an FPGA can be programmed to represent any sort of computational device. They can even be used to mimic obsolescent hardware. Another important point to note is that the power demands of FPGAs are low when compared to traditional microprocessor solutions. They have therefore been recently marketed as a “green” computing technology.

As the number of gates on an FPGA has increased in recent years, so has their applicability to general computing. This has exposed the key challenge for FPGAs as general purpose hardware accelerators – with configurability in relation to specific applications.

The traditional method by which they are programmed has been to express the required algorithm in a hardware description language such as VHDL or Verilog. This is a very powerful approach which can result in very large performance gains but suffers from a lack of programmers with the necessary skills to take advantage of the technology. To address this issue, a number of companies have been working over the past 10 years to develop higher-level language solutions – so called C-to-Gates compilers. To date great progress has been made, both with C and with the signal processing/modelling language Simulink as the source, but no single solution has yet demonstrated broad success. Many challenges also exist in simplifying and accelerating the use of the standard tools used to compile FPGA programs. At the same time, while the ability to reconfigure an FPGA during program execution exists, few examples of its use in computing have been seen to date.

Most FPGA-based solutions are currently presented as plug-in PC cards. In the future, thanks to the opening up of the Intel and AMD front-side-buses, FPGAs may be embedded directly onto PC motherboards with direct access to the PC’s microprocessor.

There therefore exists an opportunity to further develop C-to-Gates compiler solutions to solve some of the current challenges and move FPGAs further into the mainstream. An alternative approach may be to focus on the creation of pre-defined libraries (known as IP-cores) which are applicable to particular application sectors e.g. finance or bioinformatics.

GPGPUs

In 25 years the worldwide computer games console market has grown from nothing to a predicted USD46.5 billion by 2010 according to Price Waterhouse Coopers. The technology to

[17] http://www.algotronix.com

support the increasingly photorealistic games expected by gamers has evolved continuously and modern Graphics Processing Units (GPUs) deliver unprecedented levels of compute performance for this one very specific task within the computer.

It was only a matter of time before users started to consider ways in which to use the numerical processing performance of these GPUs to accelerate computationally intensive codes. GPUs used in this way have been termed General Purpose Graphical Processing Units (GPGPUs). Early experiments with these processors were disappointing because the floating-point precision of earlier generations of GPGPUs was poor.

More recently, many of the floating point implementations on GPGPUs have become IEEE 754-1985 compliant and the precision has improved greatly. Both of the leading GPU manufacturers, NVIDIA and ATI/AMD now produce GPGPUs with C-language based programming libraries for applications development. In general these processors are single-precision (32 bits) but recent announcements suggest that double-precision (64 bits) processors will be made available shortly.

Although highly powerful, GPGPUs are not a panacea. They require careful algorithm design and porting and can fail to achieve their expected potential thanks to memory bandwidth issues between the main system and the GPU. GPUs are always designed with graphics applications in mind, and other applications have to use them on a best-effort basis. There is little opportunity to influence hardware design in this area – which continues to be driven by the gaming market – but many opportunities do exist for the development of sector focussed high-performance libraries and applications based on GPGPU technology.

Other Reconfigurable Acceleration Technologies

There are many other co-processor technologies designed to take some of the computational load off the main microprocessor in a system. These include:

- Bespoke Hardware Accelerators: over the years many bespoke ASIC accelerator designs have been developed. The best known example in high-performance computing is ClearSpeed Technology's Advance range of accelerators. These provide a full 64-bit high performance accelerator and are delivered with a range of BLAS, LAPACK and other libraries commonly used in High Performance Computing.
- Digital Signal Processors: DSPs are widely used in the embedded computing and telecommunications space for signal processing. For some problems, particularly imaging and wireless baseband processing, they can provide exceptional levels of performance. In some ways their use in general purpose computing in the recent past has been eclipsed by FPGAs but for some applications they still provide the only acceleration solution.
- Physics Processors: as in the case of GPUs, the gaming sector has spawned the creation of numerical processing add-on cards dedicated to providing faster, more realistic calculations of the (mainly Newtonian) physics required by modern games. These processors can also be used as hardware accelerators in their own right for more general numerical applications.
- Application-specific ICs: large single-purpose markets, in particular in communications, often offer the volume to allow dedicated silicon accelerators to be implemented, offering lower cost and power consumption than FPGAs or DSPs. Particular examples are xDSL modems and GSM/3G handsets, which make extensive use of dedicated hardware to accelerate intensive functions. Some of these products, for example Picochip's Picoarray, are reconfigurable and could be used for other purposes, though their design often limits bandwidth or memory.

Propositions

(1) FPGA Programmability: enormous progress has been made in recent years in developing C-to-gates compilers. No one solution has yet delivered although a number of promising technologies exist. Improving the current performance of such tools will increase the uptake of FPGAs and establish a virtuous circle of FPGA usage and supported ISV application codes. Work must be done in the areas of ease of

programmability, ease of layout and standardisation of memory access and other interfaces.

(2) Numerical Libraries: to date the market for library-based acceleration on all hardware accelerator technologies has been driven by users interested in purchasing hardware from each vendor. An opportunity therefore exists to select target markets and develop library-based solutions for that market e.g. in the creative industries, the medical imaging sector or finance.

Applications

Specific functionality embedded within a larger real-time application which is used frequently and needs to be executed within a given time.

Specific functionality embedded within a larger application (e.g. a simulation), which is used frequently and enables a more fine-grain simulation with more time-steps and considering more parameters.

3.4.2 Un-met Market Needs

A successful hardware acceleration solution must either:

- address a specific functionality and deliver previously unattainable levels of performance, or
- provide a generic, easy to program solution capable of delivering reasonable performance across a wide range of applications.

In this context, the following un-met needs and requirements have been identified:

- high computational performance
- easy programmability
- fast access to data
- usage of large datasets
- linear increase of computational effort with increasing parameters.

3.4.3 Market Overview

The numerical computing market continues to grow rapidly with worldwide annual technical computing revenue forecast to grow from USD11.5 billion in 2007 to reach USD18 billion by 2012 by IDC[18].

The hardware acceleration market is dominated by hardware vendors – both large and small. The following companies are a representative sample of the sector and represent the major players.

FPGA Computing

There are two main vendors of FPGA processors:

- Xilinx Corporation[19] is the world leader in FPGA technology, with its own teams focusing on C-to-FPGA and Matlab/Simulink-to-FPGA compilers. With headquarters in California, it has a research centre based in Scotland.
- Altera Corporation[20], also with headquarters in California, is Xilinx's main competitor.

FPGA accelerator cards are offered by a number of companies. Scotland is particularly strong with respect to FPGA computing – the majority of the players being represented through the

---

[18] http://www.idc.com/getdoc.jsp?containerId=211500

[19] http://www.xilinx.com

[20] http://www.altera.com

FPGA High Performance Computing Alliance (FHPCA[21]). The worldwide leaders in this area are two Scottish companies:

- Alpha Data Parallel Systems Ltd[22], based in Edinburgh, delivers FPGA-based accelerator cards to a wide range of customers including Boeing, Rockwell Collins, JPL, Lockheed Martin, Motorola and BAE.
- Nallatech Ltd[23] is based in Glasgow and serves the defence and security industries and the HPC industry, and has recently announced a partnership with Intel Corporation.

A number of companies provide development tools (including Xilinx, Altera). C-to-Gates tool providers include:

- Celoxica[24]
- Mentor Graphics[25]
- Mitrion[26]
- Nallatech

There is clearly a big market for better tools – particularly if those tools can deliver in the general computing market.

GPGPU Computing

The major hardware vendors which supply GPGPU cards are:

- NIVDIA[27]
- ATI/AMD[28]

These cards are being used in a wide variety of projects around the world in science and business.

Other Acceleration Technologies

The leading bespoke hardware accelerator provider is currently ClearSpeed Inc[29].

There are many DSP semiconductor manufacturers including Analog Devices, ARM, DSP Group, Freescale, Infineon Technologies, STMicroelectronics and Texas Instruments. Many companies produce embedded accelerator solutions based on DSPs. Most of these vendors are now moving towards multi-core products. Companies with interesting reconfigurable hardware solutions include Icera Inc and Picochip, both with headquarters in the UK.

Two prominent leading Physics Accelerator providers have recently been bought by large corporates. AGEIA Technologies[30] has just been acquired by NVIDIA, and Havok[31] has been acquired by Intel.

---

[21] http://www.fhpca.org

[22] http://www.alpha-data.com

[23] http://www.nallatech.com

[24] http://www.celoxica.com

[25] http://www.mentorgraphics.com

[26] http://www.mitrion.com

[27] http://www.nvidia.com/cuda

[28] http://ati.amd.com/technology/streamcomputing/index.html

[29] http://www.clearspeed.com

[30] http://www.ageia.com

[31] http://www.havok.com/

Enablers

- Proven FPGA advantages such as performance and low power
- Increasing popularity of FPGA solutions

Inhibitors

In most cases two main inhibitors to their general uptake are apparent:

- Programmability: programming these devices is often a highly specialised task. Many of the solution providers have resorted to bundling numerical libraries with their devices. These library solutions can be generic or focused on a specific sector. In general, the more specific the solution the better the performance; the more general the solution the less likely that it will meet specific needs well.
- Memory bandwidth: it is very easy, with an accelerator card plugged in to a PCI-bus to saturate the bus passing data between main memory and the accelerator card. Very often this leads to much poorer than expected performance gains.

3.4.4 IP Potential

As outlined in Section 3.4.1, there are a number of opportunities for innovative technology development mainly in two categories:

- FPGA programmability: a number of companies and academic groups are already developing solutions from C or other higher-level languages to gates. However, none to date has succeeded in a general purpose way. The most successful current solutions are specific to the communications market and still require a great deal of input from hardware engineers. There are a number of potential innovations which have not been explored to date because the sector is characterised by small development teams working in isolation driven by the next customer's requirement. Standardisation could be a key enabler to progress here.
- Hardware accelerated libraries: such libraries will target a niche area where high performance is required. There are many opportunities to produce proprietary solutions which meet market needs and can be adequately protected and commercialised.

3.4.5 Timescales

The need for hardware-accelerated solutions is likely to grow over the coming years. FPGAs and less so GPGPUs are contender technologies for general-purpose hardware acceleration today and it is generally assumed they have between three to five years to deliver as general hardware acceleration solutions, before dense multi-core general-purpose processors become predominant – potentially making GPGPUs obsolete. The strength of the communications market means that FPGAs are likely to continue to progress, in terms of price/performance, for the foreseeable future. Sector specific libraries, once written, can be ported from one technology to the next.

3.4.6 Conclusions

Hardware accelerators have the potential to increase the performance of applications by faster processing and/or lower power consumption. Some specialist applications may drive a requirement for tools to ease/automate the implementation of code on hardware accelerator platforms. Monetisation of such technology will require careful consideration.

3.5 Natural Language Processing

3.5.1 Opportunity Description

Interfacing with and using machines effectively today is an increasingly complex task for humans, as rapid technological progress allows machines to become more and more complex. Typically only a small fraction of the available functionality of machines is used by humans.

Machines understanding natural language in spoken, typed or written form will provide significant advances.

The challenges stem from the highly ambiguous nature of natural languages, the large number of existing languages in the world, and the different philosophies and grammar systems of languages.

Natural language processing comprises systems that analyse, understand and generate languages that humans use naturally, so that eventually humans will be able to address machines as if addressing another person. Natural language processing is related to the disciplines of artificial intelligence[32] and computational linguistics[33].

In 1984 a very ambitious research project called Cyc[34] was initiated, which has been described as "one of the most controversial endeavours of the artificial intelligence history". Its purpose was to assemble a comprehensive ontology and database of everyday common sense knowledge, with the goal of enabling applications to perform human reasoning. Although the ultimate goal to realise a general-purpose reasoning engine has not been reached, Cycorp has provided some of its results on a no-cost licensing basis for research purposes and offers applications today in the security sector.

Call centre automation is a significant application of natural language processing technology in the form of interactive voice response (IVR) systems[35]. As communication with a call centre frequently follows typical patterns and processes, within this restricted context virtual agents are better able to understand the human side of the dialogue, analysing queries and answers, although human call centre agents are still required for more open and general dialogues like complaints. Speech analytics technology identifies customer insights, needs and wants that can be hidden in phone conversations on an ongoing basis and improves the corporate bottom line by delivering the products and services that customers want. A related area is self-service terminals, which are already commonplace in airports e.g. for check-in purposes. These terminals have the potential to be introduced in a number of other areas, from hotel reception to supermarkets and car rentals.

Research on so-called chatterbots [36]has evolved from early versions like Eliza to more sophisticated versions[37], though none have stood the Turing test[38]. In computer games, especially in role-playing genres with games like the Sims, the communication with so-called non-player characters (NPCs) is an important element of gameplay. Chatterbot technology promises to give NPCs more personality and lead to more interesting interactions for players.

Automatic language translation is another tremendously challenging target for research. Existing translation technologies like Altavista's Babelfish[39] produce only somewhat useful results, due to their inability to deal with ambiguity and cognitive concepts like humour. In practice, any documents for which (nuances of) meaning really matter, e.g. legislative or legal documents, and also product manuals or advertisements, are translated by qualified human translators[40]. While knowledge-based translation technology relies on the supply of rule systems by humans, other technologies (e.g. example-based or statistics-based) rely on raw computational power.

---

[32] http://www.aaai.org/aitopics/html/natlang.html

[33] http://www.coli.uni-saarland.de/~hansu/what_is_cl.html

[34] http://www.cyc.com/

[35] http://www.connectionsmagazine.com/articles/7/032.html

[36] http://en.wikipedia.org/wiki/Chatterbot

[37] http://www.loebner.net/Prizef/loebner-prize.html

[38] http://plato.stanford.edu/entries/turing-test/

[39] http://world.altavista.com/

[40] This sentence translated into German and then back into English by Babelfish reads: “In practice for the all possible documents (nuances of) meaning affair, e.g. legislative or legal documents, in addition, product manuals or advertisements, by qualified human translators to be really translated.” A reader can get the drift, but this is of course unacceptable for official communications.

Proposition

To develop algorithms that enable machines truly to understand natural language and communicate with humans in a natural way.

Applications

- contact centre automation (voice, email)
- self-service terminals for a range of purposes
- intelligent spam filtering
- intelligent search query execution and answer generation
- intelligent brokering (e-commerce)
- machine translation
- localisation of products, e.g. product manuals, dubbing of movies
- chatterbots (robot companion, web/application user interface)
- games, interaction with virtual characters
- advice
- semantic data mining
- automatic metadata generation
- human-machine cooperative work

3.5.2 Un-met Market Needs

The following un-met needs and requirements have been identified:

General needs:

- Human-machine interaction is increasingly difficult and inefficient, and needs to be improved in terms of quality and speed of understanding

Call centre automation:

- ability to conduct more complex and/or open-ended dialogues
- reliable speech recognition, range of languages and dialects
- adaptation to the caller
- improved TTS, prosody and range of voices
- e-mail classification / answering (fully or semi automatic)
- increasing productivity and reducing costs

Machine translation

- good (human-translator-like) quality of results
- preservation of concepts like humour
- real-time translation[41] for some applications like conferences, meetings, wiretap

3.5.3 Market Overview

Translation and interpretation services will be a USD12 billion market worldwide, according to Common Sense Advisory[42]. Large global translation companies include Lionbridge Technologies, L-3/Titan, SDL International, Systran Software Inc, and Trados Corp. Machine translation is currently a small business. It was estimated to be about USD100 million in 2007, which is about 1% of the human translation market. Most new machine translation work is funded by military and intelligence agencies, using the increasing power of computers to find statistical patterns in language.

---

[41] e.g. DARPA's Babylon programme

[42] http://www.commonsenseadvisory.com/

DMG Consulting[43] forecasts high adoption of contact centre speech analytics applications, based on past growth rates of 120% in 2006 and 100% in 2007. Speech analytics has already provided quantifiable benefits for early adopters, yielding an ROI of nine to 12 months. Frost & Sullivan estimates that the EMEA telephony-based Speech Technology and Solution Market will generate revenues for vendors of EUR455 million by 2011 from automatic speech recognition (ASR), text-to-speech (TTS) and speaker verification (SV) licences.

ComScore released their April 2007 search market share numbers showing that Google increased its market share lead to about 50%. Yahoo has about 27%, Microsoft 10%, AOL 5%, and Ask.com has 5%. Google has a stock market cap of USD150B, or USD3 billion for each 1% of search market share. Other competitors do not achieve the same revenues and market multiples, but even at the low end, 1% of market share is worth over USD1 billion. Search engine providers earn revenues of *c*. USD0.12 per search query.[44]

3.5.4 IP Potential

The understanding of human language by machines is still a big technological challenge, with IP potential in the areas of novel algorithms or the realisation of systems with massive computing power to support heuristic or brute-force approaches.

3.5.5 Timescales

The full vision of natural language processing has not materialised to date. For real progress in a platform technology sense for a range of applications, innovative algorithms are required. The futurists Ray Kurzweil predicted in 1999, that in

2009 the majority of text is created using continuous speech recognition; also ubiquitous are language user interfaces;

2009 translating telephones (speech-to-speech language translation) are commonly used for many language pairs;

2019 most interaction with computing is through gestures and two-way natural-language spoken communication.

Recent predictions of futurists are more cautious about the rate of progress:

2020 Efficient automated, wearable universal translators reduce issues of language [Wang06]

2030 Language teaching declines due to machine translation services [BT/Pearson05]

Within a 10-year horizon, it is expected that the ability of machines to understand natural language will be limited to narrow-context applications (e.g. invoice inquiries) or applications focusing on a single, very specific aspect (e.g. customer needs).

3.5.6 Conclusions

Natural language interfaces to machines able to understand speech or text could significantly enhance the utility of machines for people. The market potential is potentially large, with applications ranging from contact centre automation to Next Generation search engines. The development of such technologies represents a significant challenge, and novel solutions appear most realistic for tightly constrained applications.

---

[43] http://www.dmgconsult.com/

[44] http://dondodge.typepad.com/the_next_big_thing/2007/05/why_1_of_search.html

3.6 Cognitive Radio

3.6.1 Opportunity Description

Cognitive radio is a general term for a new kind of wireless communication system that would improve on existing systems such as GSM, 3G and WiFi. The term was coined by Joseph Mitola III in 1999 as an extension of software-defined radio (SDR).[45] A cognitive radio would seek to adapt its behaviour according to the needs of the user and the technological environment that it is operating in. The goal is to provide wireless connections with performance and a range or types of use that are not currently achievable with existing radio technologies, by taking advantage of radio spectrum that is unused in the given location.[46]

Technologies such as GSM, 3G and WiMax are centrally controlled and operated in licensed frequencies, usually by commercial operators who charge for the service or the military. WiFi, Bluetooth and similar unlicensed technologies can be used free of charge, but are allocated only very small amounts of spectrum with tight power constraints, so are subject to limitations on range and speed. In contrast, a cognitive radio would use spectrum potentially licensed for other purposes, but taking steps (for example in selecting its frequencies and power level) to avoid interference to the primary users of that spectrum.

White space radio is a subset of cognitive radio specifically targeting the ultra high frequency (UHF) radio bands that have historically been dedicated to TV. For historical and commercial reasons, TV transmissions are generally regional across the developed world, with only about one third of frequencies in use in a given location to avoid interference. UHF is particularly favourable for consumer devices because it can give good range and data rate from small, low-power devices with built-in antennas. In practice, much of the discussion around cognitive radio assumes operation in the UHF band. The term spectrum sensing cognitive radio is also used, as this type of access would necessitate the radio first detecting what other services are currently occupying the available bandwidth. In the US, the FCC has indicated that this type of cognitive radio would be authorised under certain conditions, and test are under way.[47]

Cognitive radio is sometimes claimed to offer a utopia of high speed, free wireless access without the costs and restrictions of regulation and without interfering with other users. In practice, to achieve these goals, many hurdles must be overcome. Two key technical issues, the “hidden node problem”[48] and “hole punching” (e.g. using a mediator node), are not fully understood in this context and must be addressed. Furthermore, it is known that it will be very difficult to make consumer radios that are simultaneously cheap (to compete with Bluetooth and WiFi), extremely sensitive and frequency agile.

Outside these areas, cognitive radio has other potential applications that have lower technical and regulatory risks. For example, similar systems are already under development for military use and emergency services are cited by the FCC as a key use case.

There are several connections between cognitive radio and Next Generation Computing. Cognitive radio requires higher computational performance than most radio technologies, but low power, so current developments in high-performance embedded processors is a key enabler. The “cognitive” aspect will require novel software to implement intelligence in key functions such as pairing, connection and spectrum search.

[45] http://ieeexplore.ieee.org/Xplore/login.jsp?url=/iel5/98/17080/00788210.pdf?tp=&arnumber=788210&isnumber=17080

[46] http://www.ofcom.org.uk/research/technology/research/emer_tech/sdr/ (March 2006)

http://www.ofcom.org.uk/research/technology/research/emer_tech/cograd/ (Feb 2007)

[47] http://www.fcc.gov/oet/cognitiveradio/ (Oct 2005)

[48] http://en.wikipedia.org/wiki/Hidden_terminal_problem

Proposition

Develop adaptive, two-way wireless communication technology, which opportunistically uses available unused spectrum including freed-up UHF bands, with no/minimal interference with existing users or operators.

Applications

- Consumer electronics: multimedia distribution and Internet access in the home, user-to-user (ad hoc) data connections outdoors
- Computing and mobile communications: replacement of, or complement, to WiFi, Bluetooth, UWB.
- Point-to-point and broadcast communications for emergency services, disaster relief or military use: high-availability, high-speed data links as an alternative to TETRA or other licensed services.
- Telemetry services, e.g. tracking of goods.

3.6.2 Un-met Market Needs

The following un-met needs and requirements have been identified:

- sufficient range (particularly inside buildings or in open areas) to support applications of interest
- consistently availability of high data rate bearer services to support applications of interest
- resilience to interference required for mass-market services like TV streaming and nomadic Internet access, or for emergency or military use where no other infrastructure exists.
- ease of user set-up (plug&play)

3.6.3 Market Overview

Market structure and key players

The cognitive radio market is at an early stage. Corporate proponents of the technology include Intel, Microsoft, General Dynamics, the Shared Spectrum Company and the SDR Forum.[49] The latter could evolve into the standards body that would define protocols and interoperability requirements, in a similar way to the Bluetooth SIG.

Once established, the technology could be expected to have a similar route to market as WiFi and Bluetooth – potentially bundled with those technologies on a single device – through vendors of chips and modules, for example: Intel, CSR, Broadcom, Freescale Semiconductor, Texas Instruments and STMicroelectronics.

Market size

For an emerging market such as this, size is difficult to estimate; the best figures may be obtained from adjacent markets, in particular WiFi and Bluetooth. In 2006, 213 million Wi-Fi chipsets were sold, at an annual growth rate of 32%.[50] Integrated modules (e.g. USB WiFi adaptors) are manufactured for around USD10 each. For Bluetooth, a leading vendor interviewed estimates that 1Bn units have been shipped in total by the end of 2007, with the market leader, CSR reporting revenues (predominantly from its Wi-Fi and Bluetooth products) approaching USD1billion/year during 2007. Cognitive radio could provide a significant value increment to these markets.

---

[49] http://www.sdrforum.org/

[50] http://www.instat.com/catalog/scatalogue.asp?id=3#IN0703869WT (May 2007)

Enablers

- New approaches in spectrum regulation and the freeing of TV spectrum in major markets are clearing the way for cognitive radio, but the backing of regulators such as the European Commission, Ofcom and FCC must be obtained.
- Increasing processing power and decreasing cost and power consumption of Next Generation embedded computing.
- Proof that technical hurdles can be overcome to allow cognitive radios to operate with sufficient performance while causing negligible interference, in particular to each other and to TV services.
- Addressing concerns of primary users, in particular TV broadcasters.

Barriers

- Concerns of the existing owners of the spectrum that would be targeted – especially broadcasters – must be resolved as any interference from cognitive radios would affect their customers' quality of service.
- Cellular mobile network operators present both a powerful lobby against the technology and a commercial challenge – cognitive radio would have to offer services that are much cheaper and/or significantly better than theirs, taking into account potential future cellular technologies such as WiMax, 3GPP LTE(Long Term Evolution) and femtocells.
- UHF frequencies are being offered for sale in some major markets, for example the 700MHz spectrum auctions under way in Jan/Feb 2008 in the US, and the Digital Dividend plans in the UK. Potential purchase by companies planning to offer two-way communications (in particular in the US) could create further potential competition as well as the prospect of cognitive radio having to share UHF spectrum with a wider range of technologies.

3.6.4 IP Potential

The potential for IPR creation in cognitive radio is high given the relatively early stage of the market. There is the prospect for patents in areas such as the following:

- strategies for spectrum sensing, interference avoidance, and dealing with the hole punching and hidden node issues;
- integration with Next Generation Computing platforms, other radio technologies, and antennas;
- system implementation and optimisation, to attain high sensitivity over a wide range of operating frequencies and bandwidths;
- protocols for link negotiation, integrity and interference avoidance, for example contributing to IEEE-1900-B.

IPR produced in such technologies, governed by special interest groups sych as IEEE falls into two categories: the standard and the enhancements. Patents that are adopted into the standard are subject to special restrictions on licensing and often may be included in a patent 'pool' to ensure freedom-to-operate of the early implementers.

3.6.5 Timescales

The FCC move noted above has created a market for cognitive radio in the US, though here and in other markets the main potential is expected to develop in 2012 when most developed countries complete the switch-over from analogue to digital TV – freeing up UHF spectrum and offering greater resilience to interference. In markets outside the US, significant regulatory work remains to be done to allow cognitive radio as an authorised secondary user of spectrum. Technically, a great deal of R&D remains to be done.

It is expected that research in this area to prototype level should be aiming to complete in two to three years, with commercial exploitation from about 2012. There is however significant risk that the time to market could be much longer, especially given the presence of other technologies, including UWB, and the technical and regulatory risks that cognitive radio faces: some commentators have suggested a 10-20 year timescale.[51]

3.6.6 Conclusions

Cognitive radio technology has the potential to utilise the available radio spectrum more efficiently, and could be seen as a natural progression for software defined radio technology. While regulators (and treasuries) are looking generally favourably at this opportunity, there is latent resistance from operators of existing public radio networks. Military and emergency services may be early adopters of cognitive radio technology. It is therefore expected that cognitive radio technology may be slow to emerge for mainstream, mass market, applications.

3.7 Optical Computing

3.7.1 Opportunity Description

The clock speed of today's silicon microprocessors has reached a certain limit, where further increases in clock speed are causing wasteful operation through significant power leakages. Although IBM and Georgia Tech claim that they have demonstrated the first silicon-based chip capable of operating at frequencies above 500 GHz (by cryogenically "freezing" a prototype fourth-generation SiGe chip to minus 451 degrees Fahrenheit), Bill Gates forecasted, in 2007, a 10GHz clock speed limit for practical silicon-based systems.

For today's massive server farms (e.g. as operated by Google), up to one third of budgets can be spent on electricity bills. Many-core processor systems provide one way to realise an increased system performance, but performance gains are not a linear function of the number of cores. Most applications today are programmed to run well on single processor systems, and do not scale to many-core systems without significant effort.

To improve single processor performance, new fundamental approaches need to be introduced, DNA computing and quantum computing could be such technologies, but they are still at the fundamental research stage. Optical computing appears to be a more feasible alternative technology using photons instead of electrons. Advantages are potentially much better ("light speed") performance, better signal-to-noise ratio, immunity to electromagnetic interferences, better energy efficiency and reduction of the carbon footprint, and, combined with optical communication technology, a significantly higher data throughput.

Optical computing technology could either take the place of existing silicon-based circuitry, digital as well as analogue, or lead to novel computing systems, which could be fundamentally different from classical von Neumann architectures and many-core systems. It could solve classes of complex problems which are only approached today on a heuristic basis yielding sub-optimal results. Research by Lene Hau[52] of Harvard University on "slow light" using ultra-cold atoms known as Bose-Einstein condensates (BECs) reduced the light speed to 17 metres per second (the speed of a racing bicycle) - and ultimately completely stopped a light pulse which normally travels at 186,000 miles per second. These effects could be beneficially used in future products based on optical technologies.

Fundamental building blocks include power sources like diode lasers and processing elements like photodetectors, filters, mirrors, beam splitters and modulators. As an added bonus, optical computing could increase the density of data coding beyond the two states silicon electronics can offer, much as WDM technology for fibre carries a large number of simultaneous datastreams.

An important aspect will be how the manufacturing technology required for optical computing will blend into the existing silicon processes, as chip manufacturers have invested heavily into fabrication plants.

---

[51] http://prismspectrum.blogspot.com/2006_11_01_archive.html

http://www.ofcom.org.uk/research/technology/research/emer_tech/cograd/ (Feb 2007)

[52] http://www.seas.harvard.edu/ourfaculty/profile/Lene_Hau

Proposition

To develop optical computing elements like lasers, transistor-equivalents, logic element-equivalents and further integrated circuitry, to create more powerful computing systems than can be achieved with today's silicon-based electronics technology.

Applications

- Truly optical communication networks, where all elements incl. routers and switches are working on optical principles; there are no electro-optical conversions. All-optical transponders are a key technology here.
- Switch fabric between processors/multi-cores and/or to distributed memory.
- Optical computing elements substituting electronic elements, (i) special-purpose hardware accelerators, e.g. discrete Fourier, vector matrix operations, sorting, video compression (ii) general-purpose processors
- Optical computing elements substituting analogue elements, (i) special-purpose processors capable of solving e.g. NP-complete problems in linear time, or high-end RF components, (ii) general-purpose analogue processors
- Optical computing systems with better EMV immunity, e.g. for the aerospace industry

3.7.2 Un-met Market Needs

The following un-met needs and requirements have been identified:

- increases in single processor performance
- reduction of communication lags, e.g. to support real-time applications like live broadcasts better and multiplayer games
- optimal computing solutions for complex problems, e.g. optimisation problems and simulations
- systems which are immune to EMV interferences, e.g. in aircraft and cars
- efficient usage of power and low-cost operation.

3.7.3 Market Overview

The market for carrier/enterprise routers and switches is dominated by heavyweights such as Cisco, Alcatel-Lucent and Nokia-Siemens. Infonetics Research[53] reported for the third quarter of 2007 USD3.1 billion revenues, an increase of 12% compared to the previous quarter. Worldwide optical manufacturer revenue will grow to over USD13 billion in 2010. Cisco reported USD35 billion revenues in 2007. The top two drivers continue to be 1) the ongoing migration of service providers to next generation networks based on an IP, MPLS and Ethernet data services layer, and 2) growth in consumer broadband, corporate, mobile telephony and mobile data traffic.

The computing market continues to grow rapidly with worldwide annual technical computing revenue forecast to grow from just under USD10 billion in 2006 to in excess of USD15 billion by 2011 by IDC[54]. Intel reported USD35 billion revenues in 2006.

Currently, there are no truly commercial optical computing products on the market, although products or pre-production prototypes are available for optical backplanes (e.g. Xyratex[55]), data storage (e.g. Call/Recall, Mempile) and electro-optical components (e.g. Arasor, Cyoptics, Luxtera).

Efforts leading to products in this area are likely to be costly due to the early-stage nature of much of the requisite research.

---

[53] www.info.infonetics.com

[54] http://www.idc.com/getdoc.jsp?containerId=IDC_P7

[55] http://www.xyratex.com/technology/white-papers.aspx

Lenslet[56], a company founded in 1999, intended to "design, develop and market innovative optical signal processing engines". Its EnLight Optical Processor combined optics, silicon and communication interfaces in a standard electronic board format and delivered eight Tera operations per second in 2003, *c.* approximately one thousand times faster than existing DSPs. After having raised GBP27 million in funding, the company was closed not long thereafter, with insufficient maturity of products being a major factor. Another company, Terahertz[57], spun out of Heriot Watt University in 1998, was set up with GBP3 million first-round funding plus a subsequent GBP6 million to develop optoelectronic components including planar light wave circuits, the photonic equivalent of printed circuit boards, for backplanes. Terahertz ceased to operate in late 2002, with maturity of products, cost of optical backplanes and persistence of electronics being major issues.

Routes to market can be realised through:

- Licensing of IP to vendors of high-performance CPUs (Intel, AMD, Sun, IBM, ARM), memory (Samsung, Hynix, etc), FPGAs (Altera, Xilinx), systems (Sun, Cray, IBM)
- Licensing of IP to router/switch vendors
- Potential for new fabless semiconductor firm to enter market

3.7.4 IP Potential

In the area of optical computing there is potential to generate novel IP in hybrid electro-optical circuitry by integrating photonic elements on silicon chips. There is further potential to develop basic logic elements on the basis of silicon manufacturing technology and integrate these elements into higher-order systems. In addition, there is IP potential in the development of special-purpose optical processors for specific functionality.

3.7.5 Timescales

The arrival of optical computational elements (transistors, logic gates, etc) is expected within *c*. 15 years[58]. There are still unresolved problems with cascading gates, power consumption, complex materials and laser sources, and there is still a fair amount of fundamental research needed to progress optical computing technology.

MIT's Silicon Microphotonics Technology Working Group[59] sees cost as the main driver towards photonic interconnects. They estimate that the switchover from electrical to optical systems will take 12 years for the computer serial bus and eight years for the backplane.

PA Consulting[60] expects optical computing to evolve in three stages: "The first stage will involve integrating optical interconnects between the various chips inside a computer and increasing the bandwidth between the devices. The next stage is to integrate microelectronic circuits with microphotonic circuits – the hybrid chip approach. The final stage will be everything optical."

3.7.6 Conclusions

Optical computing has the potential to revolutionise silicon-based integrated circuit technology, promising faster processing, seamless integration with fibre communications/optical bus systems and lower power consumption. Basic optical logic components have been developed in the lab, and applications of optical computing technology are expected to emerge in telecoms, datacoms and, longer-term, in hardware acceleration and general-purpose computing.

---

[56] http://www.tornado-insider.com/radar/compshow.asp?compid=4806

[57] http://findarticles.com/p/articles/mi_qn4156/is_20000924/ai_n13950745

[58] http://www.uhisrc.com/FTB/Optical/OpticalComp.ppt

[59] http://mph-roadmap.mit.edu/about_ctr/report2005/5_ctr2005_silicon.pdf

[60] http://www.paconsulting.com/services/it_services/innovation_unit/foresights/foresight_optical_computing.htm

3.8 Systems Biology Simulation

"Systems biology" is a relatively new scientific discipline which aims to gain new insights by analysing bio- and eco-*systems*, rather than studying components in isolation. The ultimate goal of a systems biological study might be, for instance, to link changes in the way an animal's body moves with biochemical changes at the molecular level within it.

The field of bioinformatics brought large-scale computation into the life sciences for the first time. Systems biology aims to extend this use of computer technology beyond data analysis to include full biochemical and biomechanical simulation of large, complex systems. This is an area of active research, and whereas many of the practitioners of bioinformatics are computer scientists, most systems biologists today have roots in the life sciences. An opportunity for new technology exists, especially based on Next Generation Computing technologies providing advanced simulation algorithms and increased processing power. Results could also be used in biologically-inspired realisations of information systems, e.g. immunity of systems to computer viruses and "mishandling" by users.

3.8.1 Opportunity Description

A good summary of the research landscape for systems biology is captured in the "In Vivo In Silico" Grand Challenge for Computer Science[61], defined in 2002 by the UK Computing Research Committee. "In Vivo In Silico" is an attempt to organise the UK computer science community to address the needs of life science. Its ultimate aim is to build a framework in which to fit sub-cellular models, a framework that can simulate models on different levels of granularity.

Starting with yeast, followed by the Nematode worm *c. elegans*, then the Arabidopsis weed the challenge is to construct computer simulations of 1,000 – 1,000,000 cells, each with a detailed intra-cellular biochemical pathway model. Single plant models can then be linked to, for instance, model fields of crops to predict yields through an ecosystem simulation. Researchers at the University of Edinburgh (Andrew Millar) expect to have this done in the virtual research lab within the next five years.

Another example is neuroscience, where scientists aim to model and simulate the human brain. Researchers at the University of Manchester (Steve Furber) have set up the Spinakker project to design a chip multiprocessor to model up to a billion neurons in real time.[62]

Systems biology is thus still firmly rooted in basic research but has clear market-relevant goals in two primary areas: human life sciences (pharmaceuticals – in silico efficacy/toxicity testing of drugs) and plant sciences (agribiotech – maximisation of yields).

Human life science – pharmaceuticals

The main opportunity identified is in the modelling and simulation of drug-relevant biological systems, e.g. organs like the heart, liver, kidney, eye, gut, reproductive system or brain. The Oxford Heart Model[63] (OHM; Denis Noble) has recently been accepted as part of the case for a new heart drug at the FDA – the first time that computer simulation has been accepted as evidence for such a hearing. The OHM is a full electromechanical simulation of a human heart, a finite element model where each element has a cell with signal pathways simulated as well. These signal pathways are the targets of heart drugs and the goal is the in silico testing of drugs. Massive computational power is needed; Oxford are currently working with Fujitsu on a petaflops (10^{15} floating point operations per second) machine for OHM.

---

[61] http://www.ukcrc.org.uk/gcresearch.pdf
[62] http://www.iee.org/OnComms/Branches/UK/England/EastAng/Camb/IETCambMar07.pdf

[63] http://www.rcuk.ac.uk/escience/news/heart.htm

Another example is the UCL Liver System Model project (David Bogle, Ann Warner), an EU Framework 7 project which builds on results of the BioBeacon project[64].

This type of cell-model-function simulation for whole organs will become more common as computational techniques are brought to bear, computing power continues to increase, and models and simulation are increasingly validated.

Agribiotech

An area that is expected to become of similar importance to pharmaceuticals in the near future is Agribiotech, particularly with regard to crop yields and biofuels.

Depletion of fossil fuel reserves, rising concerns over strategic energy security and an increased awareness of potential human impacts on climate change have resulted in recent policy shifts towards growing crops for fuel rather than food. US policy drives investment in this area e.g. through government subsidies to the farming sector, and the EU has set a target to meet 10 % of transportation fuel needs with biofuels by 2020.

The benefits of biofuels are debatable, e.g. the University of Berkeley's Tad Patzek[65] claims that corn-based biofuel production and consumption is worse for the atmosphere than burning oil. Moreover, in relation to the land area used, Patzek states that wind energy is *c*. 10 times more effective, and solar energy can be *c*. 500 times more effective than growing and converting crops into biofuels.

As the fuel market increasingly competes with food and livestock/poultry feed markets over the same crops, the price of food is expected to rise. In the US, the acreage for corn is currently increasing at 15 % per annum. The US Department of Agriculture projections for the coming years indicate the US corn supply/demand balance is changing from one of chronic surplus production to an extended period of shortages. Because of rising national demand for meat products and transport fuel, China could soon shift from being a large exporter of corn to an importer or would need to increase the rice yield from the same area of land by a factor of two by 2050, putting even more pressure on supplies.

Another alternative is algae, which can be grown using waste materials such as sewage. To supply US demand for petroleum-based fuel, a land area equivalent to *c*. 0.5% of the US would need to be converted to algae-growing facilities.

Whether for fuel or food or both, the demands on a finite area of arable land are set to increase and there is a sense that agribiotech firms are becoming interested in new simulation technologies in the same way that pharmaceuticals have been for some time. Potential applications include the improved design of fertilisers, irrigation schemes or the genetic engineering of new plant strains. As with other areas of systems biology the big challenges are in cellular pathway simulations, conquering the complexity of models and the amount of simulation required.

3.8.2 Un-met Market Needs

The opportunities outlined are of potential interest to industries like the pharmaceutical and agribiotech industry, which are very large market sectors.

The development of novel drugs today is very costly, mainly caused by a low success rate and the need for extensive clinical trials. The potential market for systems biology in pharmaceuticals lies in software tools or platforms for the multi-scale simulation of organs on large computing platforms. As illustrated by the "In Vivo In Silico" Grand Challenge and Oxford's collaboration with Fujitsu on a petaflops machine, systems biology requires a lot of computational power.

---

[64] http://grid.ucl.ac.uk/biobeacon/

[65] http://petroleum.berkeley.edu/patzek/index.htm

With a steady expected rise in world population and a looming energy crisis, there is a need to increase crop yields. In agribiotech the need is for software tools that analyse the impact of technologies on crop yield for food or biofuel or both, and suggest ecosystems which maximise the yield.

3.8.3 Market Overview

According to the report "Bioinformatics Market Update (2006)" from analysts RNCOS, the worldwide bioinformatics industry reached USD1.4 billion. This market is expected to grow at an average annual rate of 16% to hit USD3 billion by 2010. The newer market in systems biology products and services is growing even faster; according to analysts[66]. "Market projections for systems biology products and services are expected to grow at an annual compound rate of 66% to USD785 million by 2008".

In the pharmaceutical industry today, developing a drug takes an estimated 12-15 years and costs over USD800 million[67]. Up to 80% of these costs arise from failed drug candidates and computer simulation of drug efficacy is becoming an avenue of increasing interest in reducing these costs. Research suggests that even a 10% overall improvement in the attrition rate could save the industry USD242 million per drug[68].

There is currently mixed commitment to systems biology in pharmaceutical companies; Pfizer's team, for example, has doubled every year for the past four years, while others are downsizing. GlaxoSmithKline (GSK) seems to be taking an outsourcing route in their approach to systems biology, funding bio-IT start-up Beyond Genomics (BG) to "apply BG's Systems Biology technologies to GSK's ongoing study of disease and drug response in several metabolic disorders"[69]. Outsourcing is not uncommon among the big players, including Merck and Roche. Bayer are working in-house on systems biology projects and also collaborating with the likes of Physiomics plc[70], a start-up based in Oxford who provide systems biology simulation services for cancer drug studies.

Aside from in-house teams at the large pharmaceutical firms, the current major player in systems biology is IT systems and services giant IBM[71]. According to a report by IPR specialists Fernandez and Associates[72],

> *"By partnering with multiple bio-IT companies, IBM is the only company in the market that is comprehensively involved in the entire drug and discovery processes."*

While the presence of such large players is perhaps somewhat daunting, Fernandez and Associates go on to note:

> "*Taking into account the amount of work still needed to standardize data and software interoperability, along with the higher level of sophisticated software, public competition is less of a concern than IBM and other private companies. Therefore, there will be ample opportunity for bio-IT companies to compete in the systems biology market."*

Smaller niche players include the following companies:

GenoLogics Life Sciences Software Inc[73] is a Canadian company, winners of the 2005 Frost & Sullivan "Drug Discovery Technologies Niche Player of the Year" Award.

---

[5] http://www.researchandmarkets.com/reports/42935/

[67] Entelos Inc. Jan 2008.

[68] Tufts Center for the Study of Drug Development, Impact Report, Vol. 4, No. 5, 2002

[69] http://www.gsk.com/ControllerServlet?appId=4&pageId=402&newsid=133

[70] http://www.physiomics-plc.com/

[71] http://www.research.ibm.com/FunGen/

[72] http://pharmalicensing.com/public/articles/view/1117185957_4296e7a54c5ed

Entelos Inc[74] have a patented technology called PhysioLab which enables them to build "dynamic large-scale computer models of human disease", again with a focus on the integration of disparate sources of data.

Genomatica Inc[75] are "recognized industry leaders in systems biology", a San Diego-based company with a strong product line in modelling and simulation.

The agricultural side of the systems biology market is harder to assess because of its relative novelty. Systems biological methods have only recently been touted as routes to addressing issues such as global food provision and technologies are very much still in university research departments. The 2007 Gordon Research Conference on Agricultural Science included a round table session titled "How is Systems Biology Impacting Agriculture?" and featured panellists from Bayer Cropscience, Syngenta, BASF, Monsanto, Metanomics (a BASF company) and the Salk Institute[76].

3.8.4 IP Potential

A good analysis of the innovation potential for systems biology can be found in the Business Insights report of September 2005, "The Future of Systems Biology: Emerging Technologies and their Impact on Drug Discovery, Development and Diagnostics".[77].

A table of recent (1999 – 2005) patents in systems biology from Thomson Scientific is reproduced in Nature Biotechnology, volume 23, number 8[78]. The table lists nine patents, primarily to US universities.

Certainly the potential for innovation is strong; systems biology as a scientific discipline is new and commercialisation activity is in its infancy. The market seems to comprise a number of small specialist companies – undoubtedly many operating in start-up and initial product development phases at the moment – being watched by the larger pharma- and agri-biotech firms.

3.8.5 Timescales

Systems biology is a relatively new biological study field and an emerging technology area. Many technical problems still exist. Scientists and biologist are working to create, refine and retest algorithms and models to improve the quality of different approaches.

Timescales in the human pharmaceutical side are a few years ahead of the agricultural, but both are relatively new areas. Interviews with experts in this area suggest that frameworks for simulations will be completed in the laboratory over the next five years. Some examples of such software platforms already exist for specific cases but few of them have been engineered in very robust or generic ways, although there are clearly start-up companies developing technology in this area.

"In Vivo In Silico" is still a computational grand challenge. Over the next two to five years it is expected to see systems biology simulation platforms being designed and built to take advantage of increasingly parallel computing hardware.

---

[73] http://www.genologics.com/

[74] http://www.entelos.com/

[75] http://www.genomatica.com/

[76] http://www.grc.org/programs.aspx?year=2007&program=agsci

[77] http://www.researchandmarkets.com/reportinfo.asp?report_id=307455

[78] http://www.nature.com/nbt/journal/v23/n8/pdf/nbt0805-939.pdf

3.8.6 Conclusions

Systems biology simulation has the potential to improve the efficiency of drug developments dramatically, and represents a significant technological challenge. This is an emerging market with significant growth potential, and early adopters are expected to include the major pharmaceutical research and development companies.

3.9 Long-tail Software

3.9.1 Opportunity Description

Software products and services are a substantial global market and the business focus of large corporations such as IBM, EDS, CSC, Oracle, SAP and LogicaCMG. Beyond, and often in partnership with these large corporations, the market sustains a significant cottage industry of application developers offering custom solutions, e.g. for websites, e-commerce, business process automation, stock and supply chain management. However, small and medium-sized enterprises (SMEs) are not well served by mass-produced software products and services, yet the SME sector makes up over 99% of all enterprises in the EU[79]. Processes that are similar in name between businesses are actually often highly customised. The main exceptions to this are in the areas of websites and e-commerce (where around one million UK companies already have some kind of website), and customer relationship management (CRM), where Salesforce.com is proving to be highly successful.[80]

There exists an opportunity for software products and services that meet the individual needs of SMEs and workgroups of large organisations. This type of software is also called the long tail of the software market[81]. This long tail has been generally inaccessible in the past because individual custom software has been too expensive to produce, market and maintain once deployed.

Specific opportunities exist in the areas of business process automation and other productivity applications e.g. managing user-specific content and data, and in the provision of software components, which can be used rapidly to build an application to specific requirements. Typically, the only business software that SMEs currently use is accounting software from a supplier like Sage; other business processes are managed on an ad hoc basis using largely unstructured documents, spreadsheets or databases.

A related technology on the data side is BPML (business process mark-up language), a variation on UML.

Proposition

To develop software platforms for automating business processes and other productivity applications for SMEs and small workgroups within large organisations, as low-cost, managed products and services that are easy to deploy and use, integrate seamlessly with existing systems, and satisfy the individual customers' needs.

Applications

- Business processes, e.g. staffing/attendance, HR, expenses, production and stock management
- Customer management, e.g. client lists, meeting schedules, mail shots
- Supplier management
- Business information systems
- Data management systems.

---

[79] http://ec.europa.eu/enterprise/enterprise_policy/sme_definition/index_en.htm

[80] http://www.salesforce.com/uk/

[3] http://longtail.typepad.com/the_long_tail/2005/03/the_long_tail_o.html

3.9.2 Un-met Market Needs

The following un-met needs and requirements have been identified:

- customisable by any computer-literate manager, without the help of IT specialists or developers
- business process software at an affordable price point of e.g. GBP100 p.a.
- easy enough for any small business to use
- simplification or automation of common business tasks.
- high scalability is not a primary requirement – often these services will have one or at most a few users
- good fit with individual business needs
- seamless integration with mass-produced and other long-tail software products and services
- systems that are more reliable and quicker to use than the ad hoc spreadsheet-based processes that are typically implemented
- faster customisation or more effective application development based on high-level functional building blocks.

3.9.3 Market Overview

According to IDC, sales of business process management software exceeded USD1B in 2008.[82] The market for specialist business process software at the enterprise and large-scale public sector level is much larger than this, but is often dominated by one-off development projects and outsourced management contracts that can individually exceed the billion-dollar mark. An example is the ongoing NHS project to implement electronic health records for 50 million people in the UK over 10 years at a cost of up to GBP31 billion for the taxpayer.[83]

Financial data of providers of business process and accounting software to SMEs give indications of the potential market size for long-tail software:

- Salesforce.com specialises in Internet-based CRM solutions for enterprises from sole traders to large corporations, claiming 646,000 users paying from around GBP7/month upwards for the service. The company's annual turnover is about USD500 million, though it is only just reaching profit.[84]
- Sage, the provider of accounting software to SMEs, has over 5 million customers worldwide, a turnover of just over GBP1.1 billion and profits of GBP250 million in 2006/7.[85]
- Another recent example, the virally marketed Microsoft Office OneNote, has reached 10 million trial downloads and "several million" purchases at around USD99.[86]

It would therefore be reasonable to expect that the total market for long-tail software services could ultimately be worth USD250-500 million per annum.

Other companies with product or service offerings include Factonomy, a three year-old startup offering application development automation tools, Zopus, offering web front-end automation technology also used by Factonomy, and Erudine, offering modelling tools for rapid application development.

As well as the companies noted above, other companies targeting process automation software for SMEs and workgroups include Jotspot (acquired by Google in 2006), Taskserver, informu.com, Andion IT, Nsite, and IBM. Salesforce.com is also behind AppExchange, a marketplace for business productivity applications, some of which function as services, and is

---

[82] http://www.idc.com/getdoc.jsp?containerId=210144 (Dec 2007

[83] http://www.medicalnewstoday.com/articles/14829.php (Oct 2004)

[84] http://www.salesforce.com/assets/pdf/investors/Q407_Press_Release_2-21-07final_w._Financial_Stmts.pdf (Feb 2007)

[85] http://www.ar2007.sage.com/?id=11003 (2007)

[86] http://blogs.msdn.com/chris_pratley/ (August 2007)

promoting its own infrastructure (force.com) as a platform for wider business process automation applications.

A route to market might be to piggyback solutions on top of existing platform like Sage or MS Office Live.

Enablers

- Internet application server and Web2 infrastructure, and widespread broadband links between businesses and the Internet, providing platform for these services to be launched to a large market as well as a platform to deliver these services rapidly
- Growing willingness amongst users to use Internet applications such as salesforce.com
- Software must be usable and easy to customise, out of the box.

Barriers

- The market is highly fragmented and necessarily domain-specific
- Competitive risk from large players such as Google, Microsoft, IBM and salesforce.com, all with an interest in addressing this market themselves and existing projects in software as a service (SaaS)
- Lack of available components
- Lack of stable component technology standards
- Lack of certified components,

3.9.4 IP Potential

Innovation potential exists in the creation of novel software concepts and tools that allow rapid development of bespoke applications or easy customisation of standard software products and services. As with most software patents, it may take increased effort to define valid and protectable innovation in this area. Purely software innovation and business method processes are still only protectable in the US at the current time. The main competitive advantage is perceived to be time to market.

3.9.5 Timescales

This market is emerging now with the companies noted above are becoming active. Based on the activities of salesforce.com and other potential competitors, the key opportunity for R&D is thought to be in the next one to three years.

The market for software components to support long-tail services is potentially longer-term – three to five years. Although software components exist on a low level today (e.g. Java Beans) and there is a growing range of open source and proprietary components available, components with required high-level “plug & play” functionality have still not arrived.

3.9.6 Conclusions

Long-tail software products and services have the potential to provide enhanced value by being able to address specific and individual needs of businesses across a number of market sectors. The development of a platform technology flexible enough to address a large variety of needs might be challenging but, if achieved, represents a significant growth opportunity.

3.10 Market Opportunities not Prioritised

Figure 4 below lists the areas that were favourably assessed but were considered to have less potential than the seven opportunities described previously. The key reason for not pursuing these areas further at this stage is highlighted below.

Opportunity	Summary	Key reason for not prioritising
Media search engine for e-commerce, home and enterprise applications	Software algorithms to facilitate searching and matching media (images, video, audio) for e-commerce applications; search by example; search using mobile phone images	Clear unmet needs, technologies exist which address these needs to some degree (e.g. Blinx), but feasibility of technologies which could fully address these needs is unclear.
Data/memory performance optimisation of applications	Mechanisms to analyse code in terms of data processing, and transform the data processing part of code into a form which uses available processor/memory topology better, to increase the overall performance especially of data-rich applications.	For the existing hardware platforms/architectures, a number of software tools exist; For novel silicon-based hardware platforms, entry barriers appear high.
Human senses augmentation	Real time understanding of the environment to augment human senses to increase work performance and safety.	A number of prototypes have been produced in research; practical applicability and proof of efficiency in real working environment has not been demonstrated.
Real/virtual world mapping	To develop improved mapping mechanisms which capture and map aspects of the real world faster, more accurately and more frequently into virtual representations, which benefit applications e.g. in e-commerce, logistics and task planning.	Clear unmet needs. Some relatively simple functions are provided (e.g. real-time traffic information), but functions like real-time updating of maps and the built environment are not feasible (yet).
Engineering and maintenance of large heterogeneous systems	Mechanisms to engineer reliably and predictably, maintain and update large live 24/7 heterogeneous systems (e.g. NHS, tax systems).	This opportunity is part of the overarching software engineering challenge, and aspects of this opportunity are considered within the ITI SIE programme.

Figure 4: Non-prioritised market opportunities [Source: ITI Techmedia]

4 CONCLUSIONS AND NEXT STEPS

This section includes a short conclusion to the report together with the next steps that ITI Techmedia will take to investigate opportunities arising from the Next Generation Computing foresighting activity

4.1 Conclusions

Next Generation Computing touches on a number of scientific disciplines and market sectors. Not surprisingly, a large number of market opportunities have been generated in the brainstorming exercise and, indeed, this list could be extended further. This report has outlined a number of market opportunities in the area of Next Generation Computing with potential to become a future ITI programme and ITI will investigate some of these opportunities in more depth.

ITI Techmedia has reviewed the identified opportunities to assess which are the more promising area(s) to investigate further. In particular, the following criteria were considered:

- the competitive advantage delivered by meeting the unmet needs of the opportunity, i.e. assessing whether the market opportunity is credible and durable
- the key components (functional needs) which should be encompassed by a successful solution
- the potential for creating new IP
- Scottish skills base in relevant technologies
- key risks.

Following this analysis three areas have been prioritised by ITI Techmedia for further study and engagement with the membership. These are:

- Hardware acceleration
- Optical computing
- Cognitive radio.

4.2 Next Steps

ITI Techmedia welcomes R&D programme proposals and expressions of interest in related areas to assist in providing an understanding of the capabilities and commercial interest within Scotland.

If you are interested in future engagement in the area of Next Generation Computing, how to make a proposal or provide an expression of interest relevant to this area, please contact Peter Astheimer, Technology Analysis Manager, ITI Techmedia (peter.astheimer@ititechmedia.com).

APPENDIX A: GLOSSARY OF TERMS

AI	Artificial Intelligence is the intelligence of machines
Bytecode	Bytecode is a term which has been used to denote various forms of instruction sets designed for efficient execution by a software interpreter as well as being suitable for further compilation into machine code.
C/C++/C#	General purpose programming languages mainly developed and promoted by Microsoft.
CFD	Computational fluid dynamics (CFD) is one of the branches of fluid mechanics that uses numerical methods and algorithms to solve and analyse problems that involve fluid flows.
CRM	Customer relationship management (CRM) is a multifaceted process, mediated by a set of information technologies that focuses on creating two-way exchanges with customers so that firms have an intimate knowledge of their needs, wants, and buying patterns.
DRM	Digital rights management (DRM) refers to technologies which provide access control for digital media
DSP	Digital signal processor is a specialized microprocessor designed specifically for digital signal processing as opposed to analog signal processing.
ERP	Enterprise resource planning (ERP) systems attempt to integrate several data sources and processes of an organization into a unified system.
FEA	Finite element analysis (FEA) is a computer simulation technique used in engineering analysis.
Femtocells	Femtocell is a small cellular base station, typically designed for use in residential or small business environments. It connects to the service provider's network via broadband.
FPGA	A field-programmable gate array is a semiconductor device containing programmable logic components called "logic blocks", and programmable interconnects.
GPGPU	General-purpose computing on graphics processing units (GPGPU) is the technique of using a GPU, which typically handles computation only for computer graphics, to perform computation in applications traditionally handled by the CPU.
HCI	Human computer interface (HCI) is the means by which the user interacts with the computer
HDTV	High-definition television (HDTV) is a digital television broadcasting system with higher resolution than traditional television systems (NTSC, SECAM, PAL).
IPTV	IPTV (Internet Protocol Television) is a system where a digital television service is delivered using Internet Protocol over a network infrastructure, which may include delivery by a broadband connection.
JVM	A Java Virtual Machine (JVM) is a set of computer software programs and data structures which use a virtual machine model for the execution

	of other computer programs and scripts.
LTE	Long Term Evolution (LTE) is a 4th generation mobile broadband standard, the successor to GSM/UMTS
Multicore	A multi-core CPU combines two or more independent cores into a single package composed of a single integrated circuit (IC).
OpenMP	OpenMP (Open Multi-processing) is a portable, scalable model that gives programmers a simple and flexible interface for developing parallel applications for platforms ranging from the desktop to the supercomputer.
SaaS	Software as a service (SaaS) is a model of software deployment where an application is hosted as a service provided to customers across the Internet.
SDR	A Software-Defined Radio (SDR) system is a radio communication system which uses software for the modulation and demodulation of radio signals.
SoA	Service Oriented Architecture (SOA) is a software architecture where functionality is grouped around business processes and packaged as interoperable services.
SoC	System-on-a-chip or system on chip (SoC or SOC) refers to integrating all components of a computer or other electronic system into a single integrated circuit.
VHDL	VHDL (VHSIC hardware description language) is commonly used as a design-entry language for field-programmable gate arrays and application-specific integrated circuits in electronic design automation of digital circuits.
Virtualisaiton	Virtualisation in computing is the simulation of computer resources by hiding the physical characteristics of such resources from their users.
WiMax	WiMax, the Worldwide Interoperability for Microwave Access, is a wireless communication technology enabling the delivery of last mile wireless broadband access as an alternative to cable and DSL
XML	The Extensible Markup Language (XML) is a general-purpose specification for creating custom markup languages.