Chapter 2:

Agents of Change in Future Communication Systems

Hyacinth S. Nwana and Divine T. Ndumu

Applied Research & Technology Department
British Telecommunications Laboratories
Martlesham Heath, Ipswich,
Suffolk, IP5 3RE, UK.

e-mail: {hyacinth, ndumudt}@info.bt.co.uk

1 Introduction

The main goal of this chapter is to examine the roles software agents would undoubtedly play in future communications systems (FCS). First, we speculate on what FCS would look like. Next, we project that distributed software technology would be invaluable to them. We then proceed to argue the case that software agent technology would be key to providing mechanisms for realising the distributed software technology required. Drawing from work we have carried out at BT Labs, we review some state-of-the-art software agent prototypes and the roles they would play in FCS.

A word of caution — we adopt a distinctly *systems* and BT view, hence the title of this chapter; therefore, we emphasise the whole of FCS over its constituent parts as in hardware, software, services, applications, etc. Some may not agree with the views on agents in FCS which we have adopted within BT, in which case we hope this chapter provides a starting point for debate.

The breakdown of the chapter is as follows. Section 2 briefly examines some current forces, trends and lessons learned in current communications systems including in networks, telecommunications and the Internet. We do this by drawing from a diverse literature in order to posit some views about future communications systems. Section 3 emphasises the pre-eminent role that distributed software technology would play, highlighting in particular the role of agent technology. Section 4 presents the case for agent technology as a *necessary*, though not *sufficient*, component technology paradigm to realise the challenges of FCS. In Section 5, we examine briefly the current state-of-the-art regarding the realisation of FCS using several component agent solutions. In Section 6, we present what we believe society should now be doing to prepare for FCS and the agents that would run on them, emphasising the need for realistic ‘wind tunnel’ type FCS test-beds. Section 6 also outlines some of the challenges ahead for such future systems. Section 7 concludes this chapter.

2 Trends, Lessons Learned and some Viewpoints on the Future of Communications Systems

Before we proceed to discuss the potential impact of agents on communications systems, it is useful to speculate on what the latter would be like in the next century. Until recently, communications systems have been dominated by traditional telecommunications technology, i.e. telephony, telefax, etc. This is all in the process of changing and we are about to enter a new dawn of the 21st century communications systems. Drawing from a diverse literature, we posit the following about future communication systems and businesses. Of course, there is no inevitability in the evolution of any industry, not least the communications industry: it snakes and turns under pressure from many and various forces. The forces can be categorised in four classes: technological, regulatory, market structure and the needs and demands of ever growing sophisticated users. As we posit our views, we note several of these forces.

1. *Future networks will indeed be networks of networks:* Future communication systems (FCS) will consist of a wide variety of networks - fixed, optical fibre links, radio cells, optical wireless, satellite microwave, etc., all using a variety of transport mechanisms and protocols. Customers will be connected simultaneously to several different networks, each with its own limitations and advantages (Turnbull, 1994). There will be a clear demand for all these networks to inter-operate much better than they do now.
2. *Bandwidth: from dearth to glut*: Technological changes sweeping across the telecommunications industry is having two distinct effects opines Cairncross (1995b).
   Firstly, there will be glut, not the capacity shortages of the past, a point also strongly argued in Smith (1995). This glut is due to optical fibre technology which currently allows for transmission rates of up to 10Gbits/s. To put this in context, a single fibre (thinner than a grain of human hair) can carry 150,000 voice channels or thousands of compressed video channels. A single fibre can carry 30,000 simultaneous telephone conversations (Cairncross, 1995a). Furthermore, the reader may be aware that US billionaires, Craig McCaw and Bill Gates, announced their intention to construct a network in space with their Teledesic project. The project would locate and exploit hundreds of satellites in order to transmit streams of pager, voice, narrowband, broadband and video data around the globe. The goal is to create a geodesic curtain in the sky and obviate the need for increased Internet bandwidth. Though the project has met with some recent setbacks, there is no doubt that there is a growing trend in using low earth orbiting satellite (LEOS) and geo-orbiting satellites (GEOs) to provide further communications bandwidth. The Iridium phone company uses LEOs for bandwidth. Clearly, such schemes will compete to some degree with fibre technology, and with such increases in bandwidth, prices will have to tumble. As Peter Cochrane (chief research scientist and head of the Applied Research and Technology department at BT Labs) notes, bandwidth and distance have been relegated to the unimportant, and the only important parameter from a human

perspective is *delay*. He criticises the computer industry for not understanding delay and argues that any delay over 30 milliseconds destroys the user's concentration and the very process of creativity.

Secondly, the technological changes are reducing barriers to entry into the telecommunications market, and make possible new sorts of competition. Technology will also allow for huge reductions in the management of networks; maintenance currently accounts for about a quarter (i.e. 25%) of the running costs of networks.

3. *The role of the Internet in future communications systems:* A key document from the European Commission's technology ministry, headed by Commissioner Martin Bangemann (EC, 1996), spells out the rapid change the Internet is having on the increasingly communications-intensive information economy:

- 50 million people on the Internet and growing (currently) at 1 million per month.
- 48,000 networks and 2.5 million computer servers.
- 2,000 businesses currently join the Internet every month.
- 500 million PCs are predicted to be sold over the coming decade and a significant proportion of them will be hooked to the Internet.

This vision document also points out that by the year 2000, the majority of the world's financial transactions will take place online. Furthermore, 14 million buyers will spend an estimated $24 billions online. Internet Data Corp. (IDC) is also predicting an "unprecedented" growth in the use of the Internet and WWW between now and 2000; it believes this will drive the growth of electronic commerce on the web to $100 billions (as quoted in Java^{TM} Report, 1997). This report also predicts the following from a study based on over 40,000 interviews in 17 countries around the world (Java^{TM} Report, 1997):

- The number of devices accessing the Web will grow from 12.6 million world-wide at the end of 1995 to 233.3 million at end of year 2000.
- The number of users accessing the web will grow from 16.1 million to 163 million in the same time period as above.
- A third of *all* transactions are already completed over the Web, as opposed to fax or phone. By 2000, that percentage should be much larger than two-thirds.

These predictions attest to the significant role that the Internet will play in future communication systems. Indeed, it is perhaps from such predictions that George Heilmeier, CEO of Bellcore, has been quoted as saying that "the Internet will be larger than the phone system in five years" (as quoted in Clark *et al*., 1997).

Another factor that is less acknowledged about the Internet's impact on the telecommunications industry, but which is implicit in what we have already discussed above, is the sheer rate, scale and diversity of users it is attracting. The range of miscellaneous services (including education, banking, insurance, travel arrangements, health, etc.) possible on the Internet is clearly a better attraction to the largely singular telephony/fax service provided by many incumbents. Studies show that fixed telephones are only in use 1% of the time in contrast to reports that have alluded to the addictive nature of the Internet, thanks in part to chat-room and video-conferencing type services. Fransman (1998) reports on some estimates that suggest that as much as 13% of voice calls may be made over the Internet by 2002. The take-up rate of the Internet is unprecedented in technological history. As telephony gets integrated more and more with the Web such that Internet customers can speak to a sales person in order to get immediate answers at a click, customers may start learning to tolerate the poorer service currently provided by voice over IP.

Lastly, the Internet is where most users pay a flat-rate subscription, and it is obviously providing much pressure to telecommunications giants like AT&T and BT. As "more and more people communicate via the Internet for next to nothing, they will ask why telephone calls should be different" (Cairncross, 1995b). In any case, Internet telephony is already here, albeit poor at this moment in time. Cairncross' key point is that with "even a tiny share of the traffic, the Internet's impact on prices may be large".

4. *Deregulation*: Deregulation in the European Union and elsewhere is allowing a range of companies to offer telecommunications services, and is revolutionising this market. Hence, large parts of the UK national network, for example, is under the control of different suppliers; this is already the case in the US and increasingly so in mainland Europe. Indeed, as from 1st January 1998, the European telecommunications market is fully liberalised. Thus, competition is set to be fiercer in the future and BT has already lost a significant share of its business in the UK. Even in Japan, the government is trying to reduce NTT's iron grip on the local network (Cairncross, 1995a). AT&T split into three parts, the goal being to become 'slimmer and fitter' to take on its rivals. Even in countries where regulation of telecommunication players is rife, we agree with Cairncross who argues that governments can only delay the revolution - they cannot prevent it: "if they try, they will merely fail more spectacularly" (ibid). Though, ironically, much regulation has had to be introduced in order to achieve deregulation, its gains in increasing competition, reducing barriers to entry, increasing efficiency and reducing prices are already being realised.
5. *"The network is the computer"*: Sun Microsystems prophesied this some years ago, i.e. one way of viewing future networks would be as distributed computing systems. Computing and telecommunications are clearly converging, and some argue that the computer/communications revolution has already occurred (e.g. Clark *et al*., 1996). The slow but sure evolution from circuit to packet networks ensured that networks now consist of millions of small microprocessors within packet switches, blurring the difference between computing and telecommunications. Clark *et al.* correctly note that we are clearly rapidly evolving towards an information infrastructure that does not distinguish computing from communications - where everything communicates and communication is everywhere. In this regard, note that the next generation Internet Protocol (IP) is going from 32 bit to a 128-bit address space such that domestic appliances such as VCRs or camcorders may all have IP addresses.

Figure 1 illustrates such ubiquitous communications whereby users would be able to send and receive information regardless of ownership, location and format. Future communication systems (FCS), in our view, will almost certainly look like Figure 1.

There will be messaging hubs which will *seamlessly* connect virtually any subscriber device to any other back-end device, translating among formats as necessary (e.g. PDA to fax machine, phone to email, email to voice, voice to Word document, images to fax, etc.) (Reinhardt, 1994). The concept of 'Really Intelligent Networks (RINs)' (Hawker & Cochrane, 1995) will become reality as networks truly evolve into wide-area computers.

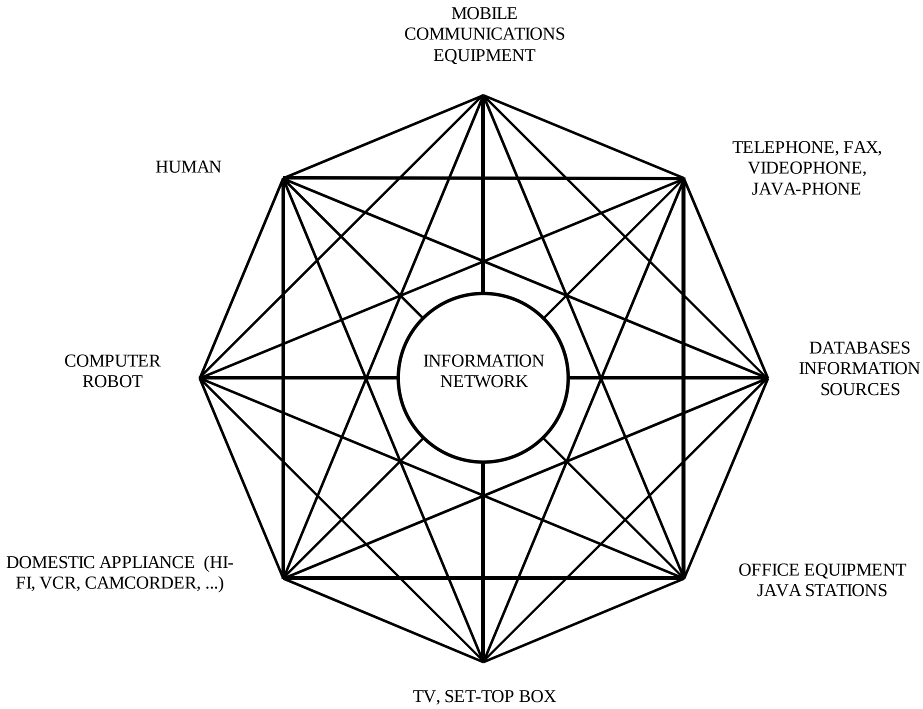

Fig. 1: Interactions between anyone and anything at anytime and from anywhere. (adapted from Hawker & Cochrane (1995))

6. *Telecommunication giants must evolve or die*: all of the above clearly point to this core conclusion. Capacity glut will ensure that the cost of a 2000 mile phone call will be equal to a 100 mile one. This is already the case in Sweden. UK Call prices across the board have already halved in real terms as compared to what they were in the 1980s. The costs of transporting bits is certain to fall, and the telecommunication giants will be saddled with expensive fixed networks. Newcomers to the market will almost certainly go wireless (as the Teledesic and Iridium projects are going), in order to avoid such costs and tackle the monopoly of the giants. There are already reputedly more mobile phones in China than fixed lines, and the trend in Poland is similar. The Internet is almost certainly going to reduce drastically telephony/fax revenues. Deregulation and liberalisation are certainly not doing the telecommunications giants any favours. In the UK alone, there are already 300+ licensed operators voraciously nibbling

away at BT's market. They would have to evolve their telecommunications infrastructures to FCS information systems (networks) and messaging hubs, as well as provide much more 'value-added' services on them in order to maintain current revenue levels. This approach of the creation of new markets via value-added networks like FCS is clearly one key long-term strategy. Current strategies to grow via globalisation, customer loyalty, getting customers to talk more, etc. are evidently much shorter-term measures.

3 Distributed Software: the Quintessential 'Glue' of Future Communication Systems

It is important to realise that a Kuhnian shift from hardware to software is occurring with the evolution to FCS as depicted in Figure 1. This is primarily because gone will be the days when the communications market is based on bits purveyance with all its associated baggage like pricing per distance. This viewpoint laid much emphasis on the provision of the necessary hardware between geographical locations to effect distributed communications.

As we have argued in Section 2, different telecommunications models will almost certainly apply in the future and companies' revenues will come from costs savings and, most importantly, from the portfolio of 'value-added' services on top of their telecommunications systems, such as electronic commerce (predicted to be a multi-billion pound industry), electronic mail, information brokering services, entertainment (e.g. video on demand), advertising, electronic distribution, information gathering/smallholding, education, digital libraries, book shops, banking, bulletin boards and communications, naturally. Our point here is that the only way to increase customer choice (in order to retain them) is to increase diversity, and FCS would need to provide for much diversity. FCS will of course have to be provably *secure* in its provision of certain services, e.g. banking.

The simple moral of the preceding paragraph is that software now becomes the pre-eminent technology of FCS, not least because customers would want to be able to have a choice of which of these services they subscribe to. In addition, they will almost certainly also want to be able to re-configure or personalise such services to suit their preferences and needs better. These are most probably only going to be realised by way of software. Furthermore, since networks are going to be distributed computer systems, they would be *programmable* in just the same way as current computers are - the only difference is that with networks the central processing unit and resources of 'this' computer are distributed, potentially across the globe. Negotiation for various services amongst service providers would be rife. This programmable capability would unleash the possibility of a vast array of possible new services, only a small fraction of which we have identified in the preceding paragraph.

Some people in the communications industry do *not* seem to share our beliefs, recognise this trend or understand this simple fact - that the future of this industry is in software, and any communications organisation ignores software at its peril. Once the First Generation FCS are in place, a positive feedback process would ensue in which customers will facilitate the identification of further services, and the companies that get to this stage first will be best placed to dominate future communications businesses. Such companies will respect no international borders and regulation as a tool (to governments) will become more and more ineffectual. The logic of every industry is that it gravitates towards where the costs are lowest, and this process will be very fluid paying little respect to international borders. Hence, if companies believe they will get a better deal from some service providers in Japan instead of the US, they will commence buying these services from Japan. Our point is that such *fluidity* would only come by way of software.

Naturally, the distributed nature of FCS (and all networks for that matter) demand the employment of software and tools of a distributed nature. *Agent technology* which we discuss in the next section is one example of such distributed software technology - one we believe will play a *necessary* though not *sufficient* role towards the provision of the sort of services we anticipate will be required of future communication systems such as those listed in Section 2.

In summary, distributed software including software agents and distributed objects will be the quintessential glue for FCS, and the full ramifications of the Kuhnian shift towards software within FCS will only become apparent with time. Imagine a smart future generation phone (which is not as dumb as that which sits on our desks today) being one of the devices sitting on FCS that provides a one-stop shop front to numerous services that would live in Figure 1. Using voice instructions, the user would be able to access services such as electronic commerce, health and education. Alternatively, the user may launch autonomous but mobile components of software – mobile software agents - onto the network to roam around and carry out specific tasks, and report the results later when they return to the user’s phone.

4 Agents in Future Communication Systems: A General Case

Agent software is arguably and currently one of the ‘hottest’ areas of software research and development. Notwithstanding the evident hype that surrounds this technology (with all its obvious dangers), the future ‘coming of age’ of this technology is (in our opinion) a virtual certainty, insofar as such predictions are accurate. Crudely, a software agent is a self-contained program capable of controlling its own decision-making and acting based on its perception of its environment, in pursuit of one or more objectives (Jennings & Wooldridge, 1996). We do not propose to overview agent technology in this chapter; the reader is referred to sources including Wooldridge & Jennings (1995), Maes (1994) and Nwana (1996).

The range of firms, universities and research institutions actively pursuing agent technology is quite broad and the list is ever growing. The scope of the applications being investigated and/or developed is arguably more impressive: currently, it really does range from the mundane (strictly speaking, not agent applications) to the moderately ‘smart’. Nwana (1996) lists examples of application domains in which agent solutions are being applied to or investigated including workflow management, network management, air-traffic control, business process re-engineering, data mining, information retrieval/management, electronic commerce, education, personal digital assistants (PDAs), e-mail, digital libraries, command and control, smart databases, scheduling/diary management, etc. Indeed, as Guilfoyle (1995) notes

"in 10 years time most new IT development will be affected, and many consumer products will contain embedded agent-based systems".

There are many reasons why the potential use of software agent technology in FCS must be investigated.

1. Software agents are subsumed by the general area of Distributed Artificial Intelligence (DAI). Thus it inherits many of the latter's motivations, goals and potential benefits. Many benefits derive from the distributed computing nature of agent technology and future communication systems. They include (Rich & Knight, 1991):

- *Natural (& geographical) problem distribution*: FCS (see Figure 1) is naturally distributed either logically or physically, e.g. its networks are distributed over vast geographical regions. Software agent technology, being naturally distributed too, maps onto such a problem domain. For example, there are different data and resources available in many geographical locations. Distributed human and software resources are already employed to handle such distributed data - why not employ distributed software agents? Assuming this, the agents may also co-operate in order to solve issues better than when they operate independently.
- *Distributed expertise:* the knowledge required to solve some problem may not reside in a single resource or agent within FCS, thereby necessitating the co-operation of many individual 'agents' in order to solve the problem.
- *System modularity*: clearly, it is preferable to maintain a *large complex* system (as would be the case with FCS-based systems) which consists of semi-independent proactive, adaptive and cooperative pieces of software (or agents), to one large monolithic one.
- *Speed*: this is a natural outcome of the parallelism and distributed nature of FCS, as agents may have their own local processor(s) and associated memories.
- *Efficiency*: clearly not all knowledge (and hence all of a large architecture like FCS) is needed for all tasks. Agent-based systems, being modular, should in principle be able to use only parts of the set-up required to solve the problem at hand.
- *Reliability/Resilience*: a multi-agent system is, in principle, more reliable because there would typically be multiple agents in the set-up which provide some particular functionality or service, i.e. redundancy would be an inherent part of FCS. Assuming this, a more reliable set-up ensues because if an agent/resource providing some functionality 'dies', another agent may take over.
- *Heterogeneous reasoning*: different parts of FCS may operate different problem-solving techniques and knowledge representation formalisms. Hence, approaches can also be combined. For example, a fault diagnosis problem would be best approached by giving fault symptoms to a number of specialist agents within FCS which concentrate on different symptoms or different areas of the system, e.g. overheating, power failure or cabling.

2. In addition to the above, software agent systems also benefit from their operation at the knowledge level (Newell, 1982). These benefits include (Huhns & Singh, 1994a):

- *Robustness*: quasi-smart agents within FCS provide for a more robust system than ‘dumb’ ones. They are also likely to degrade more gracefully.
- *Easier maintenance and reusability*: in agent-based systems generally, the communication involves high-level messages in contrast to the low-level messaging in distributed computing. This leads to lower communication costs, easier re-implementability, reusability and concurrency.
- *Platform independence*: high-level messaging facilitates the issue of platform independence, where agents may run on different platforms as long as they respond to knowledge-level messages. This is the case with many agent applications we have developed at BT Labs.

3. Software agents, in general, will provide a new paradigm for building software systems. Indeed, a new domain called *agent-based software engineering* has been invented to facilitate the inter-operation of miscellaneous software agents (Genesereth & Ketchpel, 1994). With respect to FCS in particular, we have already listed, in Section 2, some of the services that they may facilitate such as electronic commerce, electronic mail, information brokering services, entertainment (e.g. video on demand), advertising, electronic distribution, information gathering/smallholding, education, digital library services, etc. The following quotes justify to some degree these arguments:

“In the future, it [agents] is going to be the only way to search the Internet, because no matter how much better the Internet may be organised, it can’t keep pace with the growth in information…” --- Bob Johnson, analyst, Dataquest, Inc.

Negroponte (1997) similarly writes:

“I feel no imperative to manage my computer files, route my telecommunications, or filter the onslaught of mail, messages, news, and the like. I am fully prepared to delegate these tasks to agents I trust, as I tend to other matters (which could be as banal as getting dressed), while those other tasks are brought to a satisfactory conclusions. For the most part today, these agents are human. Tomorrow, they will be machines” p. 57.

Lawrence (1995) writes in the context of network management (which will apply similarly to FCS):

"By building and operating such an array of complex and interlinked systems, the telecommunications industry has set itself a big problem: how do you manage these networks?" p. 34

Management of the network may rapidly become a nightmare because centres could be swamped by management messages from this plethora of networks. We believe agents would provide one way to solve this problem.

The potential of agent technology has been much hailed, e.g. a 1994 report of Ovum’s, a UK-based market research company, is titled “Intelligent agents: the new revolution in software” (Ovum, 1994). The same firm has predicted that the market sector totals for agent software and products for USA and Europe will be worth at least \$3.9 billion by the year 2000 in contrast to an estimated 1995 figure of \$476 million (computed from figures quoted in Guilfoyle, 1995). Such predictions are perhaps overly optimistic. However, in Section 2, we provided some predictions on some of the services that FCS would provide including electronic commerce and banking. We are confident that agents would facilitate such services; hence, Guilfoyle's predictions may not be too optimistic after all.

To summarise by drawing from various strands of the preceding sections, we posit the following about agents and FCS:

- FCS will be networks of networks (consisting of fixed, optical fibre, radio cells, satellite, etc.) under the control of different service providers. Customers will subscribe to a subset of FCS networks, and the 'set union' of all the networks in FCS provides the entire FCS as depicted in Figure 1. Agents would mediate between these networks, perhaps facilitate their management, etc.
- Bandwidth and distance will be relegated to the unimportant in FCS; only delay will matter to the customers. Agents would negotiate for and provision bandwidth within the delay constraints that the customers would be prepared to tolerate or pay for.
- FCS will subsume the ever burgeoning Internet and support billions of subscriber devices (including VCRs, camcorders, etc., with their IP addresses in the new 128 bit addressing space) accessing back end devices. FCS will support the 'anywhere, any place, any time' paradigm and the translation of information between multiple modes with the aid of agents. Agents would be invaluable in locating and managing the provisioning of the distributed resources within FCS.
- FCS will increasingly support a significant fraction of the transactions that constitute the world economy in addition to thousands, if not millions, of services particularly entertainment, health and educational services. Agents already promise to play a major role in facilitating the provision of such services including electronic commerce.
- FCS will be a distributed computer system, wherein new services can be 'programmed' or configured and deployed quickly. Such services would likely make use of many other agents or services within FCS.

After this preceding 'case for agent technology' in FCS, we attempt to illustrate, albeit fragmentarily, how agents would be germane to FCS by drawing from many examples of agent research we are carrying out at BT Laboratories.

5 Agents Research at BT Labs and Future Communications Systems

Now that we have made the case for agents in future communication systems, and explored to some degree the structure and functionality of FCS, perhaps it is time to examine more closely what agents would be doing within this large complex infrastructure. In so doing, we draw from examples of agents work here at BT Labs. Generally, FCS would be a 'gold standard' deployment environment for our agent technologies because it would benefit BT enormously and give it a leading edge. So how would current state-of-the-art agents fare in such an environment? This seems to be a reasonable question to ask.

The difficulty in answering this question is that, *prima facie* at least, it assumes that the word “agent” refers to some homogeneous component of software, and many people who use the word seem to use it as such. This is a key misuse that we try to correct in Nwana (1996) by arguing that there is truly a heterogeneous body of research being carried out under this ‘agents’ banner. In the article, we present a typology of agents which includes the following: collaborative agents, interface agents, mobile agents, information/Internet agents, reactive agents, hybrid agents and heterogeneous software agent systems; see Nwana (1996) for a full explanation of this typology. Using this typology, we examine some of the functionality of current agents and how they would be deployed within FCS, in order to ‘ground’ better the discussion which we begun at the close of Section 4.

It is important to note that most current agent applications are still *demonstrators* only: converting them into real usable applications grounded in FCS may/would provide greater challenges, some of which have been anticipated but some which are currently unforeseen. However, some of these agents or variants of them are being extensively used or trialled within BT.

The example domains we highlight below illustrate better (or are suggestive of) where agents could/would be employed within FCS. A possible layered structure of agents in FCS is shown in the figure below.

User Interface Agents
Agent Interoperability/Middleware Layer
Network Services Layer
Telecommunications Network Layer

In the following table we summarise some of the prototype demonstrators we have developed and how they map to these layers and to our typology. We then proceed in the remaining part of this section to describe in some detail the prototypes we summarise in the following table, Table 1.

Agent Layer	*Some Example BT Prototypes and their genre*
User Interface Layer	Interface/Information type agents, e.g. YPA (De Roeck et al., 1998), Jasper (Davies et al., 1997), Radar, Yenta/Grapevine (Foner & Crabtree, 1997), Morse (Fisk, 1997), KSE (Davies et al., 1998), Pandora (Hyman, 1997) Collaborative type agents, e.g. PTA (Ndumu et al., 1998)
Agent/Middleware and Collaboration Layer	Collaborative-type agents, e.g. PTA (Ndumu et al., 1998), Agent-enhanced workflow (Judge et al., 1998), ADEPT (O'Brien & Wiegand, 1998), MII (Titmuss et al., 1997)
Network Services Layer	Collaborative type agents, e.g. MII (Titmuss et al., 1997), GMSN (Busuioc & Griffiths, 1996), Mobile Agents for Traffic Routing (Appleby & Stewart, 1994)
Telecommunications Network Layer	Collaborative type agents for Network management, e.g. Davison et al. (1998), GMSN (Busuioc & Griffiths, 1996) Mobile Agents for Traffic Routing (Appleby & Stewart, 1994)

Table 1 – Depicting how some of the agent prototypes we have developed map onto a layered FCS architecture

The idea of Table 1 is to help the reader have a better mental map as to how some of the prototypes we detail below map onto our posited FCS architecture. We now proceed to detail the prototypes.

It is also important to point out that this chapter does not report on the 'real usefulness' of the systems we describe. This is a key criticism of the agents area in general which we are acutely aware of, and we have begun some studies that are beginning to address this key problem.

5.1 Collaborative Agents

This category of agents emphasise autonomy and cooperation (with other agents). Hence, they may have to negotiate in order to reach mutually acceptable agreements on some matters. They typically are large coarse-grained, deliberative-style agents with non-trivial internals. With respect to future communication systems, here are some specific application areas that motivate us here at BT Laboratories:

Network traffic management: for a flavour of the size of the problem, consider the fact that BT has 5000 telephone exchanges, and in the past several years has installed

3 million kilometres of high capacity fibre-optic cable. Currently, problems such as failed exchanges are reported back to humans who handle them. This problem would be vastly greater in FCS, especially when problems occur in parts of networks that link different network providers. Network management is a naturally distributed problem so, arguably, a centralistic approach is rather unnatural. Currently, BT's telephony and derived services network are managed centrally (not including routing between switches which is carried out at the individual switch level). As networks increase in size and complexity, and different networks need to interface with each other to provide the range of services, a collaborative agent approach to management would be natural, and possibly the only way to provide a coherent method of managing the collection of networks. Introducing agents into the management of networks will yield smart networks. The interested reader should read Davison *et al.* (1998).

Workflow management: briefly, workflows consist of distinct tasks structured appropriately. Tasks refer to units of computation that perform some useful function to the system. For example, the workflow for provisioning a service that establishes a telecommunications link between New York and London can be studied. Such a service provisioning may typically take several weeks to effect, and requires co-ordination among many other operational support systems and network elements. Configuring the operational support systems may even take months. Such delays are of significant competitive significance in the business environment, and hence must be minimised as much as possible. Agents may facilitate such provisioning of services. Judge *et al.* (1998) report on some work that has been carried out in BT on agent-enhanced workflow.

Business process management/re-engineering: in corporations like BT, there is a constant need to increase the efficiency of the company, and make changes that will improve and increase on the extant services that they offer to their customers. In addition to increasing efficiency, business process re-engineering (BPR) promises to enhance organisational adaptability, flexibility and innovation in terms of their products and services. Project ADEPT (O'Brien & Wiegand, 1997), with BTL as the leading collaborator has also developed a management infrastructure which will enable BT to configure business processes dynamically using agent technology (see Figure 2).

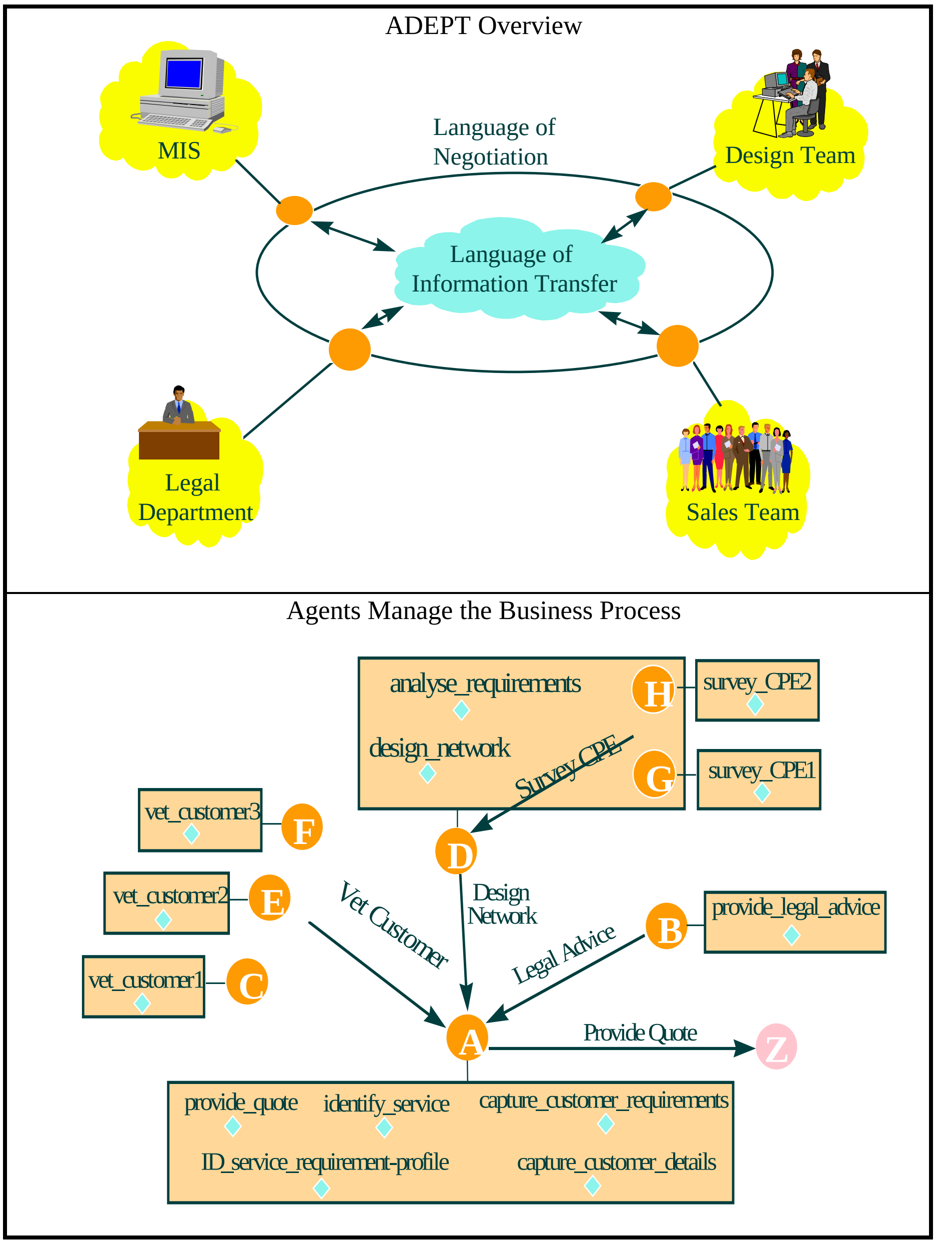

Fig. 2: (a) Concept diagram of the ADEPT business process engineering demonstrator where software agents negotiate to configure and enact a business process. (b) The nodes labeled A-H denote the agents in the scenario, each in charge of the tasks listed in its associated box.

Network fault detection, diagnosis and provisioning: fault detection involves checks performed in different regions simultaneously. This involves co-ordinating the activities of engineers separated by distance and across organisational boundaries. This problem involves distributed knowledge and control and is thus a natural domain for agents. Furthermore, separate agents managing individual networks could pave the way for rapid real-time service provisioning, and network configuration on a global scale. Individual agents could also be responsible for negotiating automatically appropriate service level agreements with the customer. The agent dealing with the customer captures the service requirements and tries to match them against the range of services on offer. Then it co-operates with the agents responsible for services and network resources, which return their best options based on their tariffs and operations costs. After receiving the options, the agent acting on behalf of its customer chooses the best match and negotiates in order to reach a service agreement. The agents could next proceed to effect the service. The BT GMSN demonstrator (Busuioc & Griffiths, 1996) demonstrated the application of agent technology to network service management and provisioning.

Personal travel assistance: imagine having your software-based personal travel assistant (PTA) that arranges trip itineraries on your behalf by integrating distributed travel services. Your PTA is all you require to provide personalised journey information and assistance. Until recently, such services would have only been provided by travel agencies, but with the Internet, such services can now be provided online. Consider, for example, your agent arranging your trip from London to New York. We have built a demonstration of just such an agent using our ZEUS toolkit (Nwana *et al.*, 1999). Figure 3 illustrates it briefly. The interested reader in this application should refer to Ndumu *et al.* (1998).

Figure 3: (see below) The Personal Travel Assistance (PTA) Demonstrator – in this demonstrator, the central agent (the PTA or travel manager) plans a trip for the user by negotiating with airline agents, hotel agents and other agents. Once agreement has been reached, the travel manager generates the trip itinerary, emails to its user and telephones the user to notify him/her that the itinerary has been generated. The envelope indicates a message being sent from one agent to another. The name server and facilitator agents are denoted ANS and Travel-Broker respectively, and the travel assistant agent is denoted Travel Manager (and shown at the centre).