OCR
RECOGNISING ACHIEVEMENT

OXFORD CAMBRIDGE AND RSA EXAMINATIONS
Advanced Extension Award

CRITICAL THINKING 9913/RB2

Reading Booklet – Documents 5, 6 and 7

SPECIMEN PAPER 3 hours

TIME 3 hours

INSTRUCTIONS TO CANDIDATES

- You will need to use Documents 5, 6 and 7 to answer the questions in Section C.

This Reading Booklet consists of 5 printed pages and 3 blank pages.

MML 3113 4/02 S35587/2

Registered Charity Number: 1066969
[Turn over

Section C

Document 5

Cyber-revolutionaries are abandoning the Web to build an anarchic, censorship-free alternative. Kurt Kleiner reports

NAPSTER's vision of free music for the masses may be dead and buried, but its spirit lives on, and not just in copycat song swapping services. When the dust settles, Napster will likely be remembered not so much for enabling music piracy as for starting a revolution that changed the way the Internet worked.

Napster is the pioneer of a technology known as peer-to-peer networking, or P2P for short. The core idea of P2P is to allow individual computers to communicate directly over the Internet. By bypassing central servers, the technology promises to transform the way people use the Net. In the process, it could destroy the ability of anyone—including corporations and governments—to control what happens in cyberspace.

"The only reason the Internet until now has been relatively censorship-free is that people who would censor the Internet haven't considered it worthwhile," says Ian Clarke, the inventor of a P2P system called Freenet. He and a group of like-minded programmers envision networks that are totally decentralised, impossible to censor and completely anonymous. In other words, cyber-anarchies.

In technical terms, P2P networks are nothing new. The Internet itself started life as a peer-to-peer system in which university and government mainframes swapped information as equals. Only when the masses began to demand access did the P2P ethos crumble. Private companies started hooking in their big computers and offering connections and online services to modest little PCs. Thus was born the client-server model. Big servers with fast connections and lots of memory hosted the information. Little computers accessed it.

Napster's winning idea was to give P2P to the masses. It figured out that it didn't have to store everything itself. Instead, it acted like a dating agency, bringing music fans—and their MP3 collections—together. Napster provided members with an index of all the music stored on other members' computers, and software that enabled them to hook into each other's hard drives. Members could then swap files without the direct involvement of Napster.

Napster was thus able to give its members access to massive amounts of music without having to store a single note itself. That turned out to be hugely popular—at the last count Napster had 61 million users—and was also a big legal advantage. It's clear that most of the recordings were being distributed in violation of copyright laws. If Napster had been storing pirated music on its site, it would have been shut down in days. The reason it lasted so long was that it could quite credibly argue that it was an innocent intermediary. If users happened to be trading pirated music it was no more Napster's fault than it's the fault of the postal service if people mail hometaped cassettes to one another.

Napster hadn't just found a way of dodging the copyright lawyers, it had solved a problem plaguing many large networks, especially the Internet. The client-server models they are built on are hierarchies, and like all hierarchies they're great as long as you are near the top. But most small-time users are near the bottom, shackled to an Internet service provider and its rules.

Among the most irksome are those rules imposed by third parties, often backed by lawsuits. Consider the successful campaign the Church of Scientology has waged against its online critics. The tactic is to accuse critics' ISPs of hosting materials copyrighted by the Church. Scientologists have taken ISPs to court and managed to have many of the critical websites removed.

Napster's Achilles' heel was that it retained a trace of the client-server model. Because members were dependent on Napster for software and indexes, record companies had a target to go after. And go after it they did. In December 1999, EMI, BMG, Sony, Warner, Universal and the Recording Industry Association of America sued Napster for copyright infringement. Although the suit is not yet settled, Napster suffered a terminal blow last month when a US court of appeal ordered it to stop enabling the exchange of copyrighted material. Napster has effectively thrown in the towel and is now trying to find a way of charging for its services so it can pay royalties.

But the P2P pirates aren't about to go away. Napster's success has inspired others, and they're determined to learn from its mistakes.

One such system is Gnutella. Originally developed by a company called Nullsoft, the software was released in March 2000 only to be withdrawn the same day under pressure from Nullsoft's parent, America Online, which was in the process of merging with music and media giant Time Warner. But the cat was out of the bag. Enthusiastic hackers unpicked Nullsoft's code and used it to write versions of their own. Within weeks there were several different but mutually compatible Gnutella knock-offs on the Web. The Gnutella commonwealth was born.

Unlike Napster, Gnutella has no central authority. No one keeps track of users and nobody indexes the files they exchange. Anyone can write software to access the network, and most of what has been written is open source, so anyone can add to it and improve on it. There are now more than a dozen versions available for free, with names like Gnotella, Newtella, Gnut, LimeWire and ToadNode.

To join the network, you simply download one of these software packages from the Web. This turns your computer into a "servent"—both a client and a server. Once you've done that you're ready to find some other servents—their locations are widely publicised on websites and chat rooms—and make contact with them. The connections are made over the Internet, and all the computers are identified by their Internet Protocol (IP) addresses, the basic numeric addresses that identify computers on the Internet. But Gnutella is not the World Wide Web. Your computer communicates directly with the servents it knows about, and those servents pass messages back and forth to yet more servents, which do the same in an ever-expanding net.

To search for a file, you type in keywords and send them to your immediate neighbours. They search the contents of their hard drives, return the hits to you and forward your request to yet more servents, which repeat the process. A single request can quickly reach thousands of computers.

Gnutella is designed to share any kind of file: images, text and software, as well as MP3s. Each user decides which files to make available. A lot of pirated material gets passed around, but the decentralised nature of the network means that there's no obvious legal target.

See you in court

The organisation most likely to start filing lawsuits is the Recording Industry Association of America. "We have not done any enforcement against Gnutella at this point. But that's not going to last long," says Frank Creighton, director of the RIAA's anti-piracy initiative. When the RIAA decides to move, he says, it will probably target that active 1 per cent. Finding out who they are shouldn't be hard because Gnutella servents need to know one another's IP addresses to communicate. Anyone can find out which ISP hosts a particular IP address, and after that a threatening letter or writ can have the user kicked off or force the ISP to reveal a name that can be pursued through the courts.

But there is a P2P network that looks capable of evading the lawyers. Called Freenet, it's a radical system created from the ground up to be anonymous and censorship-proof. Its creator, Ian Clarke, is a free-speech absolutist who feels that today's Internet, despite its freewheeling image, is vulnerable to censorship. And that's dangerous, he says. "If we look back through history we can see repeated examples where censorship and propaganda have been used to manipulate people into permitting, and even participating in, the most terrible acts of barbarism."

Like Gnutella, Freenet uses the Internet as a backbone to send and receive information, and identifies each computer by its IP address. But unlike Gnutella, it covers its tracks whenever information is transferred.

Hooking your computer up to Freenet is similar to joining Gnutella. First you download the software from the Web. Then you contact other Freenet computers, whereupon your computer becomes a Freenet "node". Freenet is made up of thousands of these nodes, and each one can make files available. When you "insert" a file—say an MP3—into Freenet it is encrypted and then copied to several other nodes. Each node knows which documents it holds and also has information about documents stored on a few other nodes. Neighbouring nodes communicate routinely, updating one another on additions to the network. But no single node knows about more than a fraction of the entire network.

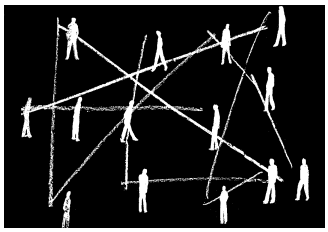

How do you get information out of a system like this? As Clarke explains it, the strategy is similar to the way people navigated before maps. Starting out, a group of travellers might have known only to go north. But the closer they got to their goal, the more detailed was the information they got from people they asked, until finally they found someone able to tell them that yes, the minstrel they were looking for lived right around the corner, second hovel on the right.

Before you start a Freenet search, you must know the title of the document you're looking for. How users will do this is still up in the air. One obvious possibility is an index within Freenet itself—though that raises the question of how to find the index in the first place. Another idea is to post it on the Web, though this may create a juicy legal target. Each document also has a numeric key that is cryptographically linked to the title, and it's this you're actually looking for during a Freenet search.

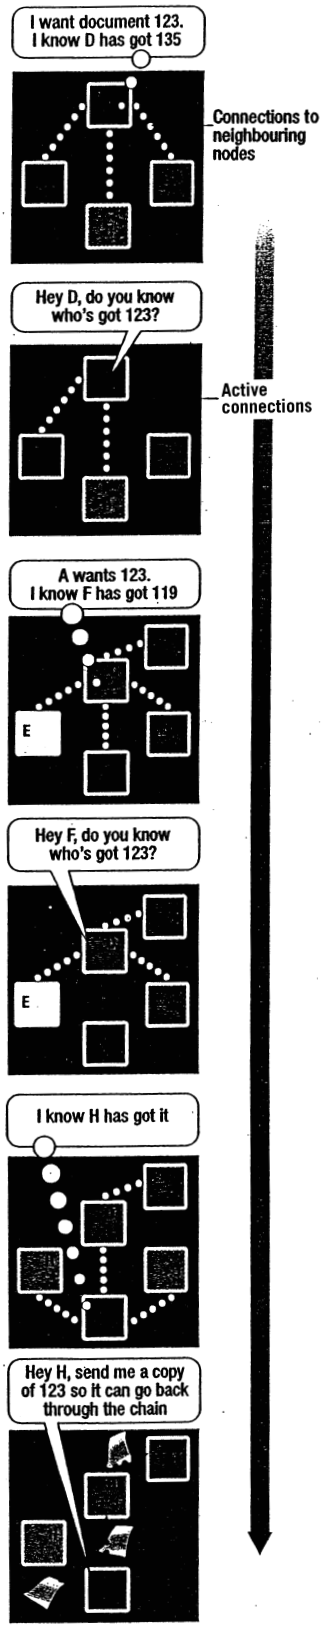

One consequence of this is that the more requests come in for a piece of information, the more copies there will be on the network, and the easier it will be to find. It also means there's no way of telling where the document originally came from. All you know is that you asked a neighbouring node for it, and it fetched the document from somewhere. Conversely, if you receive a request for a file, you have no idea who made it.

The result is a censorship-proof network. If the powers that be request a file from a node they'll get a copy. If they seize that node they'll definitely find a copy. But it would be impossible for them to prove that the file was there before they requested it, so the exercise amounts to entrapment, Clarke says. And because documents are stored in encrypted form, the node's owners can argue truthfully that they had no idea any particular document is held there. What's more, as the act of requesting a document generates new copies, censorship is self-defeating.

Let's say you know the key is 123—though of course real keys will be a lot more complex than that. Each node, including yours, knows what documents it holds, and also has a list of documents held by a few other nodes. Your computer will look to see if it has document 123. If not, it will look up to see if it knows a node that has document 123. If it doesn't, it contacts the node with the document that comes closest—maybe document 135. That node might not know where 123 is either, but it knows which node has document 119, so it sends the request there. The idea is that with each request you get closer to the document you really want. When the document is found, it's returned along the request chain (see Diagram). As the document is returned, each node along the chain makes a copy of it and stores it.

Not everyone accepts that Freenet is as censorship-proof as Clarke thinks. Creighton reckons he can bring it down by getting the IP addresses of individual nodes, sending letters to ISPs, and taking some users to court, just as he wants to do with Gnutella.

But if Clarke turns out to be correct, Freenet will usher in a different world. No one will be able to stop you downloading free music files from the Internet. You'll be able to criticise the rich and powerful without fear of being silenced or punished. And you'll be able to read whichever spy memoir your government is trying to suppress at the moment.

By the same token, you'll be powerless to stop people from plagiarising your copyrighted work or telling lies about you. Nobody will be able to take down child pornography or stolen nuclear secrets.

Napster set out to give us free music, but it seems to have put us on the road to absolute freedom of speech. If so, the real challenge hasn't even begun.

10 March 2001. New Scientist.

[Turn over

Document 6

NEW STATESMAN • 4 SEPTEMBER 1998

Sad, lonely? Log off and get out!

Of course people who surf the Internet are depressed, writes **Decca Aitkenhead**

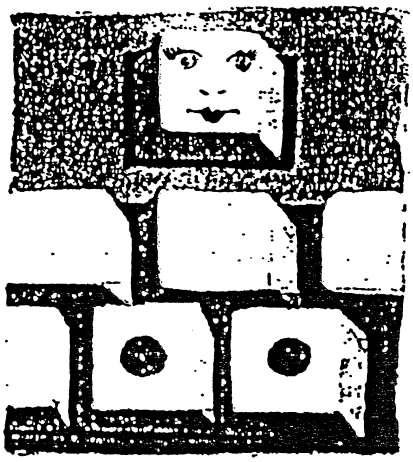

This summer, our household underwent something of a domestic overhaul. A dear American friend of our flatmate was coming to stay. This was some guy, we were assured. Boy, would he fill the room! His personality was just so, well, so *big*! We went out and bought a new double bed to accommodate his extravagant charisma, and wondered excitedly how we would keep up in his sensational company.

Our guest arrived in the middle of the night and went straight to bed. His presence only became apparent the next day when the dialling tone on the phone was found to have been replaced by a fizzle. On investigation, we discovered our guest huddled over a laptop in his room. "Whoops!" he apologised. "I'll be off in a second. I'm just on-line," he grinned.

Were it not for the fuzzy phone, it would be quite difficult now to recall that we ever had a guest to stay. He was certainly no trouble; in fact, one might recommend him as an exceptionally undemanding house guest. He was content to hunch over his laptop on the Internet all day and night, quiet as a mouse, until you found yourself forgetting he was there. Which, to all intents and purposes, he wasn't. Borrowing from his vocabulary, we might have described him as a virtual house guest.

Well, that did come as a surprise. Rather less surprising, however, was a frontpage story in the *Guardian* this week. As "new research suggests…" stories go, it was up there with "poor people have fewer bidets than company directors".

"The more that people use the Internet," we read, "the more they tend to feel depressed and lonely." Internet users have reported a decline in family interaction and size of social circle, in direct correlation with the amount of time they spent on-line. In other words, if you spend all day on the Internet you become sad and have no friends.

The golden, much-vaunted point of the internet—its defining promise—was that it would bring people together. You might appear to be alone in a Luton bedsit, but with a few taps on your keyboard you are transformed into a participant in world affairs. This is the Access All Areas scenario of your fantasies—you are driving along a superhighway, and you are welcome at every turn. Technology is delivering the promise of democracy. The globe is a village and everyone is getting to know each other.

So cyberspace is abuzz with Mike in Luton getting to know Samantha in Ohio. Mike and Samantha type out their most intimate secrets to one another in the pale half-light of the early hours. They confess dark deeds, and reassure one another. And they are led to believe that their clammily typed transatlantic relationships are real—and perhaps superior, or at least preferable, to those they have with the people asleep in the next-door bedrooms.

As it happens, Samantha's probably a truck driver called Ken. This doesn't matter, as they're never going to meet. Sadly, though, Mike isn't going to meet many real people either, because instead of leading a real life, he's stuck upstairs e-mailing "Samantha". Cyber-relations are superficially compelling precisely because they are unreal—they are the gibbering intimacies exchanged between clubbers on Ecstacy who feel they really bonded, but sadly fail to recognise one another the following day. Like Ecstasy, cyberchats feel great at the time—but leave a sickly disappointment in the stomach.

Unfortunately, there is also a tendency for Internet users to regard the web as an inexpensive form of therapy. Users off-load their miseries on-line, and anyone receiving enough of the stuff will develop a very disenchanted view on life. Internet friends are a self-selecting group of melancholics. Happier souls are, by definition outside, busy getting on with real life.

The brilliant thing about the Internet is supposed to be its lack of regulation. But if anyone is able to say anything, half of it will be lies, and so to surf the Net is to cast yourself adrift into the mad world of fantasists and lunatics. Interestingly, mad people can be made in a million different ways, but are as one in their enthusiasm for bad news. It is always the end, and never Nirvana, which is nigh. Strangely, gloomy lunatics are never so mad they can't learn how to broadcast their rubbish on the Net.

Despite all this, we keep on congratulating people for going on-line. "The Bishop of such-and-such has a website!" we squeal, as though this were some sort of achievement. Children are encouraged to swap a virtual childhood for the real thing, and so they learn nifty computer skills and have not the faintest idea how to deal with real people. But there are too many perverts in the park, their parents cluck, so they mustn't be allowed out to play.

This is the most unfortunate of ironies, since more than 80 per cent of all hits on the Internet are to pornographic sites. Like a porn mag, it offers an illusion of intimacy to disguise the isolated, addictive, desensitising reality of the experience. It is no surprise to find that it is not making people happy. The surprise is that we ever imagined it would.

Document 7

SUNDAY DECEMBER 02 2001

What kids actually do online

- Two-thirds of parents report that their children spend time talking to strangers online, and half of those parents are concerned by the sexual nature of those conversations.
- Two-thirds of parents did not know whether their children had online-only friends. The Surfsafe2001 study reveals that the majority do.
- Four in 10 parents find their children visiting inappropriate websites, and for a fifth of all parents these include adult-oriented chat and pornography.
- Placing a computer in a shared living space does not deter children from visiting inappropriate websites. Most also surf the net at a friend's home — beyond parental gaze.
- Half of the parent's express concern at what they have learned about their children's online experiences.

The good news

- The preferred interests that children under 16 pursue on the internet are, in order: **1 chat, 2 games, 3 education, 4 hobbies, 5 television & entertainment, 6 sport, 7 inappropriate topics**. Six in 10 children never seek inappropriate sites.

BLANK PAGE

BLANK PAGE

BLANK PAGE