INFORMATION HIDING
— AN ANNOTATED BIBLIOGRAPHY

There are a number of application areas in which we want to hide information or to stop someone else from doing so. These include steganography, copyright marking, the study of covert channels in operating systems, low-probability-of-intercept communications, and the study of subliminal channels in digital signature schemes.

There follow abstracts of material from number of relevant publications. A six digit number (eg. 034412) is an abstract number in 'Computer and Communications Security Reviews': see `http://www.cl.cam.ac.uk/users/rja14/` for copies of this journal.

Ross Anderson
version: 17th October 1995

'THE CODEBREAKERS'
D Kahn, *Macmillan 1967, Library of Congress catalog no. 63-16109*

The classic work on the history of crypto also contains a fair bit of material on stego. This includes a classical Chinese practice of embedding the code ideogram at a prearranged place in a dispatch; the warning the Greeks received of Xerxes' intentions via writing underneath the wax of a writing tablet; various open and jargon codes; the trick of dotting successive letters in a covertext with secret ink, due to Aeneas the Tactician; Bacon's system of encoding using two slightly different typefaces; Madame Defarge's knitting, which contained the names of enemies of the French Republic; the Cardano grill, which picks out a subset of the words on a page as being significant; and of course the whole technology of secret inks, microdots and the rest. The word 'steganography' was coined in 1499 by Trithemius, who encoded letters as religious words in such a way as to turn covert messages into apparently meaningful prayers.

People interested in policy aspects will be interested to read of the restrictions imposed by the USA in world war 2 to try and plug up as many channels as possible. The post banned a large class of objects, including chess games, crosswords, and newspaper clippings; lovers' X's were deleted; watch hands were shifted; orders for flowers could not specify either the kind of flower or the date of delivery; and items such as loose stamps and blank paper were replaced. Thousands of people were involved in reading mail, looking for language which appeared to be forced. They also rephrased telegrams; in one case, a censor changed 'father is dead' to 'father is deceased', which elicited the reply 'is father dead or deceased?'

042160 'Rechnergestützte Steganographie: Wie sie Funktioniert und warum folglich jede Reglementierung von Verschlüsselung unsinnig ist'
Steffen Möller, Andreas Pfitzmann, Ingo Stierand, *Datenschutz und Datensicherung v 18 no 6 (94) pp 318–326* (in German)

The authors describe a steganography program called DigiStilz which hides ciphertext or other material in the least significant bits of a digital (ISDN) telephone conversation. It is keyed in that the distance between two ciphertext bits is determined by a pseudorandom number, and transmission is blanked when the speech volume drops below a threshold. They tabulate test results; the rate at which the presence of ciphertext could be detected by ear varied from 1 bit in 8 for a noisy background to 1 in 64 for a quiet one. They argue that the ease with which such a system can be implemented makes legal controls on cryptography pointless.

'The Prisoners' Problem and the Subliminal Channel'
GJ Simmons, *Proceedings of CRYPTO '83, Plenum Press (1984) pp 51–67*

In order to get round the US government's prohibition on publishing work on steganography, Simmons introduced the following abstract model: Alice and Bob are prisoners, and wish to hatch an escape plan. All their communications pass through the warden, Willy. If Willy sees any ciphertext in their messages, he will frustrate them by putting them into solitary confinement. Simmons shows that a message authentication without secrecy channel providing r bits of message authentication can be perverted to allow an $l < r$ bit covert channel between the transmitter and a designated receiver at the expense of reducing the message authentication capability to $r - l$ bits. Under quite reasonable conditions, the detection of even the existence of this covert channel can be made as difficult as the underlying cryptoalgorithm. In view of this open — but indetectable — existence, the covert channel was called the "subliminal" channel.

'The Subliminal Channel and Digital Signatures'
GJ Simmons, *Advances in Cryptology—EUROCRYPT '84, Springer LNCS v 209*

The author shows how subliminal channels can be implemented in the ElGamal and Schnorr signature schemes.

034412 'Subliminal Channels; Past and Present'
GJ Simmons, *European Transactions on Telecommunications v 5 no 4 (Jul/Aug 94) pp 459–473*

This paper describes a protocol proposed in 1978 to monitor compliance with the SALT II treaty, and how the author found a potentially disastrous flaw in it — a subliminal channel which would have allowed the USSR to discover which Minuteman silos contained missiles. This discovery led to his subsequent work on the topic of subliminal channels, some of which is described. In particular channels in the ElGamal and DSS digital signature scheme are compared, and it is shown that the DSS provides the most hospitable setting for subliminal communications discovered to date.

'How to Insure that Data Acquired to Verify Treaty Compliance are Trustworthy'
GJ Simmons, *Proceedings of the IEEE v 76 (1984) p 5*
This is one of a series of papers in which the author describes the evolution at Sandia National Laboratories of a solution to the problem of verifying compliance with a nuclear test ban treaty. Data from sensors placed on another country's territory must report certain types of information but not others.

023619 'Subliminal Communication is Easy Using the DSA'
GJ Simmons, *Eurocrypt 93 pp 218–232*

The author continues the discussion of subliminal channels in the digital signature algorithm. He focusses in this paper on the broadband channel - the one which requires the recipient to know the sender's secret key, and uses examples to compare it with the similar channel in the standard El-Gamal signature scheme. He concludes that DSA provides the best subliminal channels so far discovered.

'Protocols for Data Security'
R DeMillo, M Merritt, *IEEE Computer v 16 no 2 (Feb 1983) pp 39–50*

In one of the earliest papers on cryptographic protocol failure, the authors pointed out that a protocol for playing poker over the telephone leaked information via quadratic characters.

034110 'Electronic Marking and Identification Techniques to Discourage Document Copying'
J Brassil, S Low, N Maxemchuk, L O'Gorman, *IEEE Infocom 94 pp 1278–1287*

The authors describe an AT & T system to discourage illegal document copying by marking documents with line-shift encoding. Successive lines of text are shifted up or down by $1/300''$, thereby encoding a serial number. Experiments on postscript documents showed that even third photocopies could be scanned and decoded with very few errors; different detection techniques are discussed, as are the effects of image defects and the countermeasures available to infringers.

034157 'Electronic Document Distribution'
NF Maxemchuk, *AT & T Technical Journal v 73 no 5 (Sep/Oct 94) pp 73–80*

This article describes techniques developed by Bell labs to make it easier to identify people who redistribute electronic documents. Identifying marks can be encoded in line spacing, word spacing and font features; their relative resistance to photocopying and to various erasure strategies is discussed. Such techniques can be used, for example, to embed a purchaser's name and credit card number in a document. A trial run is scheduled for an issue of the IEEE Journal on Selected Areas in Communications in 1995.

042120 'Copyright Protection for Electronic Publishing Over Computer Networks'
AK Choudhury, NF Maxemchuk, S Paul, HG Schultzrinne, *IEEE Network v 9 no 3 (May/June 95) pp 12–20*

The authors discuss the security requirements for electronic publishing; the reasonable goal might be to create a system which is as secure as paper, in that large-scale low-cost anonymous redistribution would be prevented. One solution would be to put trusted hardware in printers; where this is not possible, decryption software is the obvious choice and reverse engineering the obvious threat. They discuss some techniques which might make this less likely or less feasible.

'Towards Robust and Hidden Image Copyright Labeling'
E Koch, J Zhao, *Proceedings of 1995 IEEE Workshop on Nonlinear Signal and Image Processing (Neos Marmaras, Halkidiki, Greece, June 20–22, 1995)* and
`http://www.igd.fhg.de/www/igd-a8/pub/IEEE_Hidden.ps`

The authors describe a system developed to hide copyright messages in JPEG compressed digital images. These messages are robust in the fact of lossy compression and low pass filtering; they are encoded at a number of locations in the picture which are determined by a secret key.

'Embedding Robust Labels Into Images For Copyright Protection'
J Zhao, E Koch, *Proc. Int. Congr. on IPR for Specialized Information, Knowledge and New Technologies (Vienna, Austria, August 21-25, 1995)* and
`http://www.igd.fhg.de/www/igd-a8/pub/EmbedLabel.ps`

The authors describe two copyright marking schemes — one for embedding labels in JPEG compressed images, and one for black and white images. The former alters the quantisation table in the discrete cosine transform, and the latter encodes a bit according to whether a block of pixels is more black than white. Varying the table or block size allows the tradeoff between robustness and visibility to be selected.

'Video-Steganography: How to Secretly Embed a Signature in a Picture'
K Matsui, K Tanaka, *Proceedings, Technological Strategies for Protecting Intellectual Property in the Networked Multimedia Environment, Journal of the Interactive Multimedia Association Intellectual Property Project, v 1 no 1 (Jan 1994) pp 187–205*

The authors discuss schemes to embed copyright messages in pictures compressed using a number of different signal processing schemes, including predictive coding, ordered dithering, facsimile and discrete cosine transform. For example, fax messages are signed by replacing the rightmost pixel of each line with the next bit of the signature.

'Assuring Ownership Rights on Digital Images'
G Caronni, *preprint, Institut für Technische Informatik und Kommunikationsnetze, EHT Zürich, 3/94*

The author develops a technique for marking digital images by dividing the image into blocks and encoding successive bits as the mean value of each block. The marks can be made harder to detect by varying the size and position of the blocks in the image, which is needed anyway to avoid regions where altering the brightness would be obvious. This technique is robust against distrotion and lossy compression.

'Watermarking Digital Images for Copyright Protection'
FM Boland, JJK Ó Ruanaidh, C Dautzenberg, *Proceedings, IEE International Conference on Image Processing and its Applications, Edinbrugh 1995*

The authors propose to watermark an image by dividing it into blocks and modifying the coefficients of a transform of each block. This transform can be a discrete cosine, wavelet or even Walsh transfrom, depending on the losses and distortions to which the encoded picture may be subjected. They measure the amount of signal that can be added to each region of the picture — in highly textured blocks, energy is more evenly distributed.

042155 'Speakeasy: The Military Software Radio'
RJ Lackey, DW Upmal, *IEEE Communications Magazine v 33 no 5 (May 95) pp 56–61*

The authors describe a US Army project to develop a military radio which uses programmable signal processing to emulate more than 15 existing radios, including GPS, cellular phones, satcom, analogue HF, SINCGARS, HAVE QUICK and low probability of intercept modes. It will operate from 2MHz to 2GHz with interchangeable RF modules and IF digitisation; in addition to software waveform processing, there is also a programmable security processor which emulates the comsec and transec features of five existing crypto devices.

'Schutz der Vertraulichkeit des Aufenthaltsorts von Mobilfunkteilnehmern'
H Federrath, J Thees, *Datenschutz und Datensicherheit (June 1995) pp 338–348 (in German)*

The authors describe how direct sequence spread spectrum techniques can be used to make it harder for an eavesdropper to determine the phyical location of a transmitter. Low-probability-of-intercept systems use cryptographic rather than maximal length spreading sequences; a direction finding apparatus would need to do a autocorrelation or crosscorrelation analysis of the sequence in order to select the signal. By comparison, frequency agile systems are relatively easy to intercept. There are a number of other engineering factors which need to be considered when designing anonymous networks, such as computer security measures to protect location information stored in the network, and dummy traffic to preventing active attacks through the network. Call setup is troublesome, particularly when initiated by the mobile station; the most elegant solution involves Rabin's beacons.

'Cryptology and the origins of spread spectrum'
D Kahn, *IEEE Spectrum v 21 no 9 (Sep 1984) pp 70–80*

This article describes SIGSALY, the first digital secure telephone, which was used by Roosevelt and Churchill during the war. It used a vocoder with 10 bands of 300 Hz, each sampled for amplitude every 20mS; the digital signal was Vernam encrypted (though since the samples had six levels, the arithmetic was modulo 6). The cables from the scrambling equipment to the users were pressurised and alarmed. Finally, the radio link used an early spread spectrum technique to reduce the likelihood of interception or jamming. One of the inventors of spread specturm was the actress Hedy Lamarr, who obtained a US patent in 1941 on a frequency agile torpedo control system.

'Adaptive Fuzzy Frequency Hopper'

PJ Pacini, B Kosko, *IEEE Transactions on Communications v 43 no 6 (6/95) pp 2111–2117*

The authors describe a system which uses fuzzy logic to adaptively generate a pseudorandom sequence for frequency hopping radio. It is trained on a sample pseudorandom sequence and is claimed to be harder to intercept than a traditional shift register sequence.

'SPREAD SPECTRUM SYSTEMS WITH COMMERCIAL APPLICATIONS'

RC Dixon, *Wiley 1994, ISBN 0-471-59342-7*

Spread spectrum techniques provide a number of benefits, most notably signal hiding, jamming margin, selective addressing, multiple access, interference rejection and high-resolution ranging. For nineteen years, Dixon was the standard reference work on spread spectrum, but was relatively unknown outside a circle of military specialists. Recent advances in integration, as well as a 1985 decision by the FCC to allow commercial use, have led to rapid growth in commercial applications ranging from GPS to digital cellular telephones; Dixon has just rewritten his book to take account of all this.

The mechanics of spread spectrum systems can be quite complicated. The basic techniques — direct sequence, frequency hopping and time hopping — are simple enough in concept, but there are many complex tradeoffs between error rate, process gain, chip or hop rate, synchronisation, and the various strategies available to participants and opponents; and the linkage between coding, cryptographic and RF engineering aspects is uniquely complicated. Dixon provides a guide to the underlying theory which should be accessible to a graduate student, and goes on to discuss the engineering aspects of satellite uplinks, GPS, military tactical radios and modems, digital cellular radio and vehicle location. There are also hundreds of references which provide a lead into the research literature.

'METEOR BURST COMMUNICATIONS: THEORY AND PRACTICE'

DL Schilling (editor), *Wiley 1993, ISBN 0-471-52212-0*

The ionised trails of micrometeors entering the earth's atmosphere reflect radio waves, especially in the low VHF band, and attempts have been made since the 1960's to use this phenomenon for communication. Although the initial interest in the subject faded with the introduction of satellites for the bulk of beyond-line-of-sight communications, meteor burst communications are now used in a number of commercial and military rôles. Their intermittent nature, and the relatively small ground footprint of each trail, make them inherently hard to monitor.

This book is the first comprehensive guide to the subject to appear in modern times. Such a guide is welcome, as the subject spans a very wide range of subject matter: from the physics of the meteor trails themselves through the various coding and other techniques which are used to maximise the available channel capacity through to a number of issues arising from practical engineering experience. The context of the book is a US Air Force network in Alaska, which provides backup communications for early warning radars in the event that satellite communications are knocked out.

024130 'Laser communications for covert links'

JL Jaeger, RT Carlson, *Laser Communications 93 pp 95 - 106*

The authors describe a prototype laser communications system built by MITRE and others to replace military line-of-sight RF communications. The devices operate in the infrared; this not only increases the difficulty of detection, but also provides eye safety at all distances. The design is described, and extensive field tests reported:

depending on the weather, the link operates at 1,024, 128 or 16Kbps, and at the lowest of these rates, the system can cope with all conditions except for heavy fog or rain. Excellent covertness is claimed.

034127 'Tracing Traitors'

A Fiat, *Crypto 94 pp 257–270*

In satellite TV applications, a rogue user might leak either the key or the plaintext of an encrypted broadcast to a pirate audience; the operator will try to track such users by giving each user a personal decryption key with which he can decrypt a session key from an enabling block; thus pirate decoders can be traced. However, an attacker might take several subscriptions and try to generate an innocuous personal key from them; the author discusses various combinatorial schemes which can be used to prevent this.

'Anonymous wireless rings'

K Diks, E Kranakis, A Malinowski, A Pelc, *Theoretical Computer Science v 145 (10/7/95) pp 95–109*

The authors intriduce anonymous wireless rings; each processor has a single buffer, and cannot tell which of its neighbours the arriving bit has come from. This increases anonymity over the conventional wired ring, but turns out to be computationally equivalent, both interms of the kind of functions which can be computed and the bit complexity of input collection. A suit of bit collection algorithms is presented for both synchronous and asynchronous operation. Open problems include symmetry breaking and leader election.

'Preserving Privacy in a Network of Mobile Computers'

DA Cooper, KP Birman, *Oakland 95 pp 26–38*

The authors tackled the problem of location security of mobile computers by developing a replicated memory service which allows users to read from memory without disclosing which locations they are reading, and without wasting too much bandwidth. Their solution uses a MIX-type network which bundles up read requests, and a store which uses secret sharing techniques to spread each word over a number of locations. Read and write protocols are described.

'A Cautionary Note on Image Downgrading'

C Kurak, J McHugh, *Proceedings of the 8th Annual Computer Security Applications Conference pp 153–159*

The authors study how one image can be embedded in another; their concern is identifying the threats posed by downgrading classified satellite images. They give examples of pictures of an aircraft, an airfield, and text, embedded in each other by replacing the four least significant bits of each pixel of the cover picture by the four most significant bits of the embedded picture. The loaded pictures cannot be distinguished by eye, and the quality of the recovered hidden pictures is acceptable for intelligence purposes. They conclude that manual inspection alone is insufficient to prevent image downgrading being used to leak information.

'An EMACS based downgrader for the SAT'

J McHugh, *Proceedings of the 8th NCSC; reprinted in 'Computer and Network Security' pp 228–237*

The author describes a system he developed to enable US military personnel review classified information for downgrading. It was inspired by emacs, but targeted at a Honeywell A1 machine and formally verified. It leads the officer through the text a paragraph at a time; when he has cleared it, it is reformatted. Both surplus whitespace and non-printing characters are removed.

'A Note on the Confinement Problem'

B Lampson, *Communications of the ACM v 16 no 10*

This paper pointed out the existence of covert channels. These channels arise where a resourse is shared between two entities between which we wish to inhibit communication; for example, a virus attached to a 'SECRET' process might signal classified information down to another virus at 'UNCLASSIFIED' by modulating the CPU load or the position of a disk head.

'A GUIDE TO UNDERSTANDING COVERT CHANNEL ANALYSIS OF TRUSTED SYSTEMS'
NCSC-TG-030 version 1, November 1993

This is the official NSA guide to the identification and elimination of covert channels in multilevel secure systems. Military multilevel secure systems — at least at the higher levels of evaluation — should limit covert channel bandwidth to about one bit per second. The techniques involved include both channel elimination and noise insertion.

'Covert Channel Capacity'
J Millen, *1987 IEEE Symposium on Security and Privacy*

This is one of the classic papers on calculating the capacity of a covert channel using entropy equations; it shows how information-theoretic properties can also be represented in automata-theoretic terms.

'The Influence of Delay on an Idealized Channel's Bandwidth'
IS Moskowitz, AR Miller, *1992 IEEE Symposium on Security and Privacy pp 63–67*

The authors analyse the relationship between the bandwidth of a covert channel and the underlying queueing parameters.

041205 'An Entropy Conservation Law for Testing the Completeness of Covert Channel Analysys'
R Browne, *Fairfax 94 pp 270–281*

The author defines a complete set of covert channels as one which can operate to produce the maximum covert information flow. He shows that such sets are characterised by their satisfying an entropy conservation law, in that a fully informed onlooker perceives an output uncertainty equal to the covert capacity plus the relevant noise. This in turn lets the system behaviour be expressed in a kind of normal form.

'Extended Abstract: An Architecture for Covert Channel Control in Real-Time Networks and MultiProcessors'
R Browne, *Oakland 95 pp 155–168*

The author introduces an architecture for dealing with covert channels in multilevel realtime networks and multiprocessor systems. He proposes a 'systolic, elastically separable' system; systolic means phase-delayed I/O and non-preemptive scheduling, while elastically separable means that each multilevel task can be rewritten as a fixed, finite sequence os single level tasks which read down but do not write up, and such that only tasks at the highest level in the system refer to a realtime clock. He shows that such systems confine covert communication in useful ways.

021136 'Architectural Implications of Covert Channels'
N Proctor, P Neumann, *Proc 15th NCSC pp 28 - 43*
This paper reviews covert channels: how they occur, what assumptions are needed to ignore them, how to eliminate them from resource allocation algorithms and what the tradeoffs are. It then proposes an architecture for eliminating them and describes a design for a multi-level disk drive using manual allocation. This drive can allow read-down and write-up operations which have no covert channel but still yield adequate performance. The authors argue that building secure operating systems is beyond today's technology, and argue that using single-level processors with a multi-level disk gives maximum assurance at a reasonable cost.

021229 'The Channel Capacity of a Certain Noisy Timing Channel'
IS Moskovitz, AR Miller, *IEEE Transactions on Information Theory v IT-38 no 4*

(1992) pp 1339 - 43

A covert timing channel may suffer noise generated by time sharing delays as other users compete for resources. Two strategies for communicating in the presence of this noise are analysed and the resulting channel capacity is determined.

034217 'Covert Channels — Here to Stay?'

IS Moskowitz, MH Kang, *Compass 94 pp 235–243*

Capacity is not the only measure of a covert channel, as channels may exist which have a one-off capability to send a fixed number of bits and thus possess zero capacity. One should rather measure the longest message which can be transmitted under realistic assumptions. There are a number of trade-offs in minimising this, and various mechanisms including pumps may be useful.

032819 'Simple Timing Channels'

IS Moskowitz, AR Miller, *Oakland 94 pp 56 - 64*

Shannon's definition of channel capacity as $C = lim_{n\to\infty}\frac{\log|S_n|}{n}$ (where the S_n are sequences) should actually be $C = limsup_{n\to\infty}\frac{\log|S_n|}{n}$, as the ordinary limit does not exist in many cases of practical interest, especially in the analysis of timing channels in multilevel secure systems. A new proof of the capacity bound is given based on the use of z-transforms; this reduces capacity bounds to radii of convergence, and can be used to determine the capacity of a number of complex timing channels. These typically turn out to be the roots of a real trinomial of the form $1-x^{-a}-x^{-d}$ or $1-x^{-a}-x^{-a-d}$, which can be found in closed form using special functions.

031218 'A Pump for Rapid, Reliable, Secure Communications'

MH Kang, IS Moskowitz, *Fairfax 93 pp 118 - 129*

Existing techniques for dealing with covert channels in multilevel systems include blind write-ups and periodic read-downs. The authors propose instead a data pump with both low and high buffers, and discuss its capacity; this can be improved by judicious use of randomisation, and in practice there should be no performance penalty in benign situations.

'A Network Version of the Pump'

MH Kang, IS Moskowitz, DC Lee, *Oakland 95 pp 144–154*

The authors show how to extend the NRL data pump from a single system into a distributed one. The basic pump keeps a buffer and emits ACKs to Low at probabilistic times based on a moving average of the ACK times from high; this is extended to multiple buffers at both Low and High levels. One must pay attention to congestion control issues; if inputs could send messages faster than outputs could handle them, then covert channels would open up.

021229 'The Channel Capacity of a Certain Noisy Timing Channel'

IS Moskovitz, AR Miller,*IEEE Trans. on Information Theory v IT-38 no 4 (1992) pp 1339 - 43*

A covert timing channel may suffer noise generated by time sharing delays as other users compete for resources. Two strategies for communicating in the presence of this noise are analysed and the resulting channel capacity is determined.

021420 'A Classical Automata Approach to Noninterference Type Problems'

IS Moskowitz, OL Costich, *Proc Franconia 92 pp pp 2 - 8*

The authors are concerned with covert channels in a computer system that supports concurrent processes of differing security levels or clearances. They assert that a 'Secure Nondeterministic Automaton' is useful for identifying unauthorised probabilistic channels.

023217 'Modelling a Fuzzy Time System'

JT Trostle, *Proc Oakland 93 pp 82 - 89*

The fuzzy time technique used by DEC in the VAX security kernel randomises the period between clock ticks in order to minimise the capacity of the bus-contention covert channel. However, a scheduling channel is constructed which bypasses this countermeasure; a model is developed which can be extended to multiple hosts; and finally more secure scheduling algorithms are discussed.

042183 'Performance Analysis of a Method for High Level Prevention of Traffic Analysis Using Measurements from a Campus Network'
BR Venkataraman, RE Newman-Wolfe, *Computer Security Applications 94 pp 288–297*

The authors investigated the cost of rerouting and padding the University of Florida's network traffic to provide resistance to traffic analysis. They considered the leakage through a covert channel where a '1' was signified by the presence of traffic on the link from node i to node j. Making the traffic spatially neutral thus corresponds to blocking traffic analysis; thus traffic was rerouted and padded to put the same amount of traffic on each link. Doing this perfectly is a hard linear programming problem, so a number of heuristic solutions were tried and measured.

'Capacity Estimation and Auditability of Network Covert Channels'
BR Venkatraman, RE Newman-Wolfe, *Oakland 95 pp 186–198*

The authors continue their research into how the type, amount, and timeliness of data traffic on a network can be used to send covert information. They discuss their work in light of Browne's ideas of mode security. An audit threshold is analyzed with respect to the maximal damage due to the above types of covert channels. A study of the University of Florida network shoed that attempts to monitor traffic patterns in detail generated unmanageable amounts of data. Further details can be found in the first author's dissertation.

'Untraceable electronic mail, return addresses, and digital pseudonyms'
D Chaum, *Communications of the ACM v 24 no 2 (Feb 1981) pp 84–88*

In this classic article, the author introduces mix-nets (anonymous remailers). These decrypt incoming traffic, add or remove padding, reencrypt it and dispatch it in lexicographically ordered batches. Mechanisms are also discussed for anonymised return addresses, digital pseudonyms, blinded certified mail, and the use of a hierarchy of subnets to provide scalability. The possible application discussed is digital elections.

'Covert Channels in LANs'
CG Girling, *IEEE Transactions on Software Engineering v 13 no 2 (Feb 87) pp 292–296*

The author reports an experiement to transmit covert information across a network by choosing the destination of the next packet out of a set of addresses. Different LAN speeds and block sizes were tried, and the resulting performance measured; the capacity is $(\log_2 K . \log_2 N)/M$ where there are K patterns, N interfaces and M time units between patterns. Finally, he discusses the possible roles of cryptography and traffic padding in reducing the bandwidth.

'Networks Without User Observability—Design Options'
A Pfitzmann, M Waidner, *Advances in Cryptology—EUROCRYPT '85, Springer LNCS 219*

In normal communication networks, operators and intruders can easily observe when, how much and with whom the users communicate, even if the users employ end-to-end encryption. Once ISDN is used for almost everything, this could becomes a severe threat. There are, however, a number of technical options to keep the recipient and sender (or at least their relationship) unobservable; the authors consider some possible implementations and extensions, and propose some performance and reliability enhancements.

'The Dining Cryptographers Problem: Unconditional Sender and Recipient

Untraceability
D Chaum, *Journal of Cryptology v 1 (1988) pp 65–75*

Keeping confidential who sends which messages, in a world where any physical transmission can be traced to its origin, seems impossible. The solution presented here is based on passing messages round a ring of participants; it is unconditionally or cryptographically secure, depending on whether it is based on one-time-use keys or on public keys, respectively. It can be adapted to address efficiently a wide variety of practical considerations.

'Security without Identification: Transaction Systems to Make Big Brother Obsolete
D Chaum, *Communications of the ACM v 28 no 10 (Oct 1985)*

By partitioning consumer information into separate unlinkable domains through the use of user-created "digital pseudonyms," the dangers inherent in large-scale automated transaction systems, as currently structured, can be avoided.

'How to Break the Direct RSA-Implementation of MIXes'
B Pfitzmann, A Pfitzmann, *Advances in Cryptology—EUROCRYPT '89, Springer LNCS 434*

MIXes are a kind of anonymous remailer, suggested by David Chaum in 1981. If RSA is used as this cryptosystem directly, i.e. without hashing to destroy the multiplicative structure, the resulting MIXes can be broken by an active attack which is perfectly feasible in a typical environment. The attack does not affect the basic idea of MIXes, provided they are implemented carefully; but it does show that present security notions for public key cryptosystems may not suffice for a system which is to provide a service such as anonymity. We also warn of attacks on further possible implementations of MIXes, and we mention several implementations which are not broken by any attack we know.

'Unconditional Sender and Recipient Untraceability in Spite of Active Attacks'
M Waidner, *Advances in Cryptology—EUROCRYPT '89, Springer LNCS 434*

A protocol is described to send and receive messages anonymously using an arbitrary communication network; it is unconditionally secure. This improves a result by Chaum: The DC-net guarantees the same, but on the assumption of a reliable broadcast network. Since unconditionally secure Byzantine Agreement cannot be achieved, such a reliable broadcast network cannot be realized by algorithmic means. The solution proposed here, the DC^+-net, uses the DC-net, but replaces the reliable broadcast network by a fail-stop one. By choosing the keys necessary for the DC-net dependently on the previously broadcast messages, the fail-stop broadcast can be achieved unconditionally secure and without increasing the complexity of the DC-net significantly, using an arbitrary communication network.

032627 'Breaking an Efficient Anonymous Channel'
B Pfitzmann, *Eurocrypt 94 pp 339 - 348*

At Eurocrypt'93, Park, Itoh, and Kurosawa presented two efficient designs for an anonymous channel based on Chaum's mix-nets. The idea was to simulate a trusted host for applications like electronic voting with secret inputs and public outputs. Here, the author first identifies a passive attack against both designs that may allow correlation of inputs and outputs, and shows how to avoid this by careful choice of parameters. She then demonstrates an active attack that completely breaks one of the designs; there may be ways to avoid this attack on the other design by adding counters and redundancy.

'Receipt-Free Mix-Type Voting Schemes'
K Sako, J Kilian, *Eurocrypt 95 pp 393–403*

The authors use Chaum's anonymous channel to construct an anonymous voting scheme under which voters cannot prove how they voted and so cannot so easily be bribed or coerced. However each voter can still check that every vote was properly counted.

'Covert Distributed processing with Computer Viruses'
SR White, *Advances in Cryptology — Crypto 89, Springer LNCS v 435 pp 616–619*

Viruses can be used to perform distributed processing without the knowledge or consent of the machine owners. The class of problems for which such processing might be useful is discussed, and includes keysearch; a virus which infected ten million machines might break a DES key in about three months.

'A SHORT COURSE IN COMPUTER VIRUSES - SECOND EDITION'
Frederick B Cohen, *Wiley 1994 ISBN 0-471-00769-2*

Fred Cohen is uniquely qualified to write about computer viruses, having invented them twelve years ago as his PhD project. In addition to a brief history of the subject, including covert channel attacks on military secure Unix systems, he ranges from theoretical results on undecidability through to practical accounts of the latest virus and anti-virus techniques.

034404 'Recursive mappings for computer virus'
XA Li, JH Fu, YG Song, HY Yang, *Chinacrypt 94 pp 279–286 (in Chinese)*

This paper builds on Adleman's suggested definition of computer viruses in terms of Gödel numberings, and suggests a recursive mapping definition. They show that programs exist which are viruses under their definition but not Adleman's and vice versa; and that for any virus, there exists a program which it cannot infect or injure.

'Reflections on Trusting Trust'
K Thomson, *Communications of the ACM v 27 no 8 (Aug 84)*

In his Turing award lecture, Ken Thomson shows that a Trojan horse can be hidden in a compiler; it inserts copies of itself every time either the compiler or the login program is compiled.

032134 'Visual Cryptography'
M Naor, A Shamir, *Eurocrypt 94 pp 1 - 12*

The authors show that visual information can be protected in such a way that it can be decoded using the human eye rather than by a computer. They hide pictures by splitting them into seemingly random patterns of dots: when these are superimposed, the picture appears. The technique has perfect secrecy (in Shannon's sense), and can be used to hide signatures on ID cards so that they are only legible through a special filter. It can also be generalised to a k out of n secret sharing scheme (for small k and n).

'Mimic Functions'
P Wayner, *Cryptologia v XVI no 3 (July 1992) pp 193–214*

The author shows how the inverse of Huffman coding can be used to endow one file with the statistical characteristics of another. Examples are given of randomly generated text with the statistical properties of English, for third order through sixth order statistics and for a context free grammar.

'Strong Theoretical Steganography'
P Wayner, *Cryptologia v XIX no 3 (July 95) pp 285–299*

The author proposes a system of steganography based on context free grammars. Learning a van Wijngaarden grammar is equivalent to a Turing machine. In a pratical system, input ciphertext bits can be used to select productions from the grammar, and can be recovered through parsing so long as the grammar is unambiguous. An example is given based on English text generated using these ideas.

'The Intel 80x86 processor Architecture: Pitfalls for Secure Systems'
O Sibert, PA Porras, R Lindell, *Oakland 95 pp 211–222*

Traditionally, compusec evaluations have paid little attention to hardware. However most products evaluated B2 or above use Intel processors, and this motivated the authors to look for potentially exploitable hardware features. They found covert channels in the Task Switched flag, segment attributes, interrupt descriptor tables, debug registers, the time stamp counter, and performance counters. Caches can also provide covert channels, and in addition to all these design features there have been a number of implementation errors reported, inlcuding freak security flaws. There are also a number of subtle traps for TCB implementers. Thorough analysis is made more difficult by the fact that the processors are proprietary.