Cribs

Question 1.

a)

•The service area is covered by approximately hexagonal cells to achieve tight packing (tessellation). In a practical system the actual cells will be designed to overlap to some extent and their actual shape will depend upon local topography

•A cellular system comprises a network of Base Stations (BSs), which communicate with mobile stations (MSs)

•An MS communicates with the BS serving the cell in which it currently resides

•When the MS moves from one cell to another, 'handoff' or 'handover' from the current BS to the BS serving the new cell occurs

•High capacity is achieved by:

–**Dynamic channel allocation** – i.e., channels are only allocated to users when they are required.

–**Frequency reuse** – A systematic way of using the same channels over and over again. This is necessary since only a finite number of channels are allocated to an operator. In this case the same channels are reused in cells separated by a specified minimum distance. This distance (known as the *reuse distance*) is governed by the maximum level of interference that can be tolerated from nearby cells (*cochannel* cells) without causing unacceptable quality of service.

b)

•To increase capacity, what options are available?

–Allocate more frequencies N – difficult owing to congested spectrum

–Go digital to reduce p

–Employ directional antennas (known as sectorisation) at the base stations to reduce interference and hence reduce the effective reuse pattern

–Reduce the size of cells (known as microcells) so that fewer users occupy each cell – can be expensive owing to the requirement for more base stations to be deployed

•Essentially the allocated spectrum is divided between clusters having p cells, e.g., 3, 4, and 7

•As p increases, so does the distance between cochannel cells:

–This improves the received SNR

–At the cost of reducing capacity, since the total number of channels N, must now be divided amongst a larger number of cells, p. That is a maximum of N/p channels (calls) can be allocated per cell. For reasonable user satisfaction, the probability of reaching this level of calls (blocking) should be quite low, say < 5%.

•The minimum value of p that can be used depends upon the minimum SNR that can be tolerated at a receiver.

•Old analogue systems (e.g., TACS) require a min of 18dB SNR, requiring the use of $p = 7$. Digital modulation schemes using ECC (as in GSM) can tolerate much lower SNR values and so can use smaller reuse patterns, e.g., $p = 3$ or 4.

•Clearly this raises system capacity

c)
•Interleaving and Frequency Hopping combined with Error Correction Coding are the main techniques used to combat interference and fading channels.
•Two types of interleaving are employed:
–Bit interleaving: Essentially shuffles bits within one speech frame
–Block interleaving: Moves whole block of shuffled bits into different transmitted bursts, giving further protection against bursty errors resulting from signal fading. The aim of interleaving is to break-up bursts of errors due to signal fading causing them to appear as single errors at the input to the error correction decoder, thus improving its performance
•Frequency Hopping
–217 hops/s keeps max fade duration down to around 4.6 ms.
–When combined with half rate ECC (convolutional) and interleaving over 8 hops gives immunity to most types of fading
•When combined with FH, the interleaver is capable of breaking-up the bursts of errors even when the mobile is stationary
–Also randomises cochannel interference, thus minimising size of the reuse pattern (which would otherwise have to be designed for the worst case interferer), thereby raising system capacity

Question 2
a)
"Network-centric" or "server based" computer and communications
Concentrate applications and services on centralised managed facilities
Protocols to deliver only user-interface over networks to simple end-user equipment
Avoid user-data or un-recoverable state on the end-user equipment
Terminals simpler, complete system easier to manage

b)
Size – less hardware, smaller and more appealing footprint
Environment – no disk or fan, less heat and noise pollution
Longevity – simpler design will last longer, computing resources upgraded at centre
Stability – terminal is more stable, application set can be managed better centrally
Mobility – can move desktops around between terminals
Security – terminal is no longer a weak point, concentrate security efforts at centre
Cost – lower total cost of ownership

c)
X Window System
High level (applications are built in knowledge that graphics may be remote)

Xclient process runs on networked server machine
Xclient is application code (eg Xterm, Xclock) linked against X libraries

Xlibraries are programmatic interface to X protocol

X protocol is device independent operations eg create, move, delete window
Window is any element such as text, label, pixmap, hierarchically organised
X protocol provides network transparency (endianism, socket can be local/remote)

X server process runs on terminal equipment
X server decodes X protocol, builds data structures
Maps operations on to actual devices (screen, keyboard, mouse)
Accepts connections from multiple X clients

Plus: Open source
Plus: Still alive (Linux, KDE, GNOME, OpenOffice)
Minus: Unix centric
Minus: Bad at network interruption/terminal reboots – Xclients die
Minus: Terminal is stateful, Xclients rely on state built up in Xserver
Minus: No session mobility
Minus: No session sharing
Minus: Poor compression, bad over slow links

Microsoft Terminal Server
Middle level (applications don't know graphics are remote, but operating system does)

Standard Windows applications and operating system front end
Use standard Windows Messages (keys/mouse) and GDI (graphics display interface)

Terminal server provides virtual device interface for each user
Virtual keyboard, mouse and graphics driver backend
Virtual devices generate RDP (remote display protocol)

Win Terminal is RDP client
Terminal decodes RDP and maps on to actual keyboard, mouse and graphics drivers.
Terminals made by a number of manufacturers

Plus: Efficient (works at the OS driver level)
Plus: Good compression, reasonable over slow links
Plus: Supported by Microsoft
Plus: Session mobility
Minus: Microsoft centric, not good for mixed environments
Minus: Terminal Server license expensive
Minus: Proprietary protocol

VNC (Virtual Network Computing)
Low level (neither applications or operating systems know graphics are remoted)

VNC server is application level agent on server end
Snoops public OS interfaces, generates VNC protocol

VNC client is software running on terminal
Decodes VNC protocol and maps on to actual keyboard, mouse and graphics

Plus: Open source
Plus: Highly portable, and thus very interoperable
Plus: Session mobility

Plus: Session sharing
Plus: Good stability (doesn't replace OS or drivers)
Minus: Less efficient than RDP (have to work hard to get hints from OS)
Minus: Not multi-user on windows

SunRay
Solaris Servers run SunRay server software
SunRay server software comprises GNOME desktops and StarOffice applications
Generates low level VNC like protocol

Ship over proprietary sub/ip level private network

Proprietary terminal decodes protocol and maps onto keyboard, mouse, graphics

Minus: Completely closed world, private network island
Minus: Proprietary hardware terminal
Plus: SmartCard reader on terminal for session logon and security
Plus: Hot desking
Plus: Sharing

d)
Consider all costs of procuring, installing and maintaining a system, including personnel
Consider over a period of time, eg yearly average over five years

Direct costs (equipment) – terminals, servers, networking equipment
Direct costs (facilities) – power, heat, noise, air conditioning
Direct costs (software) – software procurement, license management
Direct costs (personnel) – system management, installation, upgrades, support
Indirect costs - self help and peer-peer help, inefficiency, frustration

Question 3
a)
i.
TESLA was designed to solve the problem of source authentication in the following context.
One source periodically broadcasts packets to many recipients.
Recipients will experience different delays and may occasionally even lose packets.
Recipients cannot afford to verify a public key signature on each packet.
The scheme should be efficient.

ii.
TESLA builds upon the earlier "Guy Fawkes" protocol but adds some improvements.

Each packet P_i is authenticated by a MAC. The sender divulges the MAC key only after the packet and its MAC have been received and can no longer be tampered with.

To prevent a man-in-the-middle opponent from
- forging a packet,
- then computing a valid MAC with an invented key,

- and then releasing the invented key,
the keys themselves must be validated. In TESLA, they are part of a Lamport hash chain. To validate a key k_j, the verifier only needs a previously-issued known-good key k_i from the same hash chain. Verification consists of hashing k_j for j-i times and checking whether the result is k_i.

Since the verification described in the above paragraph only requires "any previous key" as opposed to "the immediately previous key", the verification can still be carried out even after the recipient (= verifier) experienced packet loss.

The MAC key k_i that allows verification of packet P_i is only divulged after P_i has been sent. But *how long* after P_i was sent? If too soon, a man-in-the-middle attacker might receive k_i before a recipient receives P_i; if this possibility exists, the recipient should not trust P_i and source authentication will be impossible. If too late, the recipient experiences a long delay between receiving the packet and being able to verify it; in some time-sensitive cases (multimedia, stock tickers etc), the packet may become "stale" and therefore useless by the time its source has been validated. To fix this problem, in TESLA each packet is protected by several MACs. The keys to these MACs are released at different time intervals. Recipients who enjoy low network latency can use the early keys for verification and enjoy timely authentication; while recipients who suffer significant delays are still able to verify their packets with later keys.

iii.
The fields of packet P_i will be:

Payload M_i. Carries the information the sender wants to send.

Multiple MACs of this payload with different keys: MAC(K_i_1, M_i), MAC(K_i_2, M_i), ..., MAC(K_i_m, M_i). The first index says which packet this key protects. The second index says which Lamport chain this key belongs to.

Multiple revealed keys: K_j1_1, K_j2_2, ..., K_jm_m. Indices as above. The keys are relative to already-transmitted packets, so all the jx indices are < i.

b)
see pictures in handout

Question 4
a)
UWB or impulse radio is the use of extremely short duration pulses (sub-nanosecond) instead of continuous waves to transmit information.
The pulse directly generates a very wide instantaneous bandwidth signal according to the time scaling properties of the Fourier transform relationship between time and frequency.
Occupied Bandwidth >> Information Bandwidth.
It has a very low duty cycle 1/1000 or less.
UWB fractional bandwidth = $(fh - fl) / fc > 25\%$ or total BW > 1.5GHz.
It has low power consumption and a low probability of detection signature.

b)

The system consists of active tags which chirp out UWB pulses and receivers which are placed within a building.
Position is calculated from multiple sensor readings.
Multilateration is used to calculate the 3D position (ie time difference of arrival).
Four readings are required to calculate 3D position.
More readings provide better accuracy (outliers are thrown out).
Non-linear regression can be used to refine position.
Accuracy depends on the level of time synchronisation and the position of receivers.
Actual performance of systems ranges from 10cm to 100cm 95% of the time.

c)
Orientation can be obtained by having more than one tag attached to a rigid object
It may also be possible to obtain orientation by observing which sensors received the signals.
However this may be problematic because this technique requires the body to shadow the signal which happens with ultrasound but probably much less true with UWB.
Received signal strength might also be useful.

d)
In order to make the system scaleable the location system is made cellular.
An additional control radio is used to trigger individual tags.
If the UWB system combines communications with location a second radio system is not required.
When a tag moves from one cell to another a handover of the control link takes place to another base station.
The control frequencies can be reused in the normal way.
Additional sensors attached to the can be used to help determine which tag should be prioritised (eg trembler switch).
Higher level analysis (eg tag is moving) can also be used to schedule tag transmissions.
Applications should indicate their level of interest in particular objects.