Management Information Systems
Rutgers Business School / Undergraduate New Brunswick
Professor Eckstein, Spring 2007

Class Notes

Class 1 – Overview, Course Rules

Overview

- Topic: using computer and network technology to help run businesses and other organizations
- Won't focus especially on "managers"
- Will combine "Top-down" descriptive learning (the RTP book) with "bottom-up" learning by example (Microsoft Access and GB book)

Rules and Procedures – see the syllabus and schedule

Class 1 – Basic Definitions and Concepts (Chapter 1 of RTP)

Data, Information, and Knowledge:

- *Datum* is singular, *data* is plural (book says "data item" and "data items")
- *Information* is data structured and organized to be useful in making a decision or performing some task
- *Knowledge* implies "understanding" of information
    - Example from book: company analyzes its recruiting data and concludes that recruits from school *X* tend to have good outcomes only if their GPA's are at least 3.0. In future, based on this "knowledge", they screen applicants from school *X* by their GPA's, only interviewing those with at least a 3.0 GPA.
    - One common kind of knowledge representation in computers is called "artificial intelligence" (AI). It got a lot of hype in the 1980's, and then went somewhat out of fashion, but it is still growing gradually. We will not discuss it much, and stick to "information" instead.

Information systems (definition of some basic terms)

- The ways that organizations
    - Store
    - Move
    - Organize
    - Manipulate/process

  their information
- Components that implement information systems – in other words, *Information Technology*
    - Hardware – physical tools: computer and network hardware, but also low-tech things like pens and paper

    - Software – (changeable) instructions for the hardware
    - People
    - Procedures – instructions for the people
    - Data/databases
- Information systems existed before computers and networks – they just used relatively simple hardware that usually didn't need software (at least as we know it today). Strictly speaking, this course is about "CBIS" (Computer Based Information Systems). Because of the present ubiquity of such systems, we usually leave the "CB" to be implicit.
- Impact of computer and network hardware and related software/services (Table 1.1):
    - Can perform numerical computations and other data processing much more quickly, accurately, and cheaply than people
    - Can communicate very quickly and accurately
    - Can store large amounts of information quickly and cheaply; retrieval can often be very rapid
    - Can automate tasks and processes that previously required human labor (various degrees possible, of course)
    - Information doesn't have to be "stuck" with particular things, locations, or people
    - Can have a downside
        - Small glitches can have much wider impact (minor software bug grounds all aircraft in Japan)
        - Fewer people in the organization understand exactly how information is processed
        - Sometimes malfunctions may go unnoticed (American Airlines yield management story)
- *Information architecture* is the particular way an organization has arranged its information systems: for example, a particular network of computers running particular software supports the marketing organization, while another network of computers running different software supports the production facilities, etc.
- *Information infrastructure* consists of the hardware and software that support the information architecture, plus the personnel and services dedicated primarily to maintaining and developing that hardware and software.
- *Application* and *application program* are somewhat fuzzy terms, but typically denote computer software and databases supporting a particular task or group of tasks.
    - Example from book: HR uses one application to screen job applicants and another to monitor employee turnover
    - A classic business IT problem: applications that don't communicate with one another (effectively)

Class 2 – Types of information systems

Refer to RTP Figure 1.2:

- *Departmental information systems*, or *functional area information systems* are designed to be operated within a single traditional functional department of an organization such as sales, human resources, or accounting. In the early days of CBIS, these were often the only kind of systems that were practical.

- *ERP (Enterprise Resource Planning) systems* are a relatively extreme reaction to the problem of poorly integrated functional area systems, offered by vendors such as SAP, Oracle, and PeopleSoft. They aim to support the entire organization's needs with essentially a single integrated system. They have enormous potential benefits, but are also notoriously tricky and expensive to configure and install.
- *Transaction Processing Systems (TPS)* gather data about everyday business events in "real time" as they occur. Examples:
    - You buy 3 items at a local store
    - A shipment of coffee beans arrives at a local distribution center
    - A package is unloaded from a FedEx or UPS aircraft

  All of these events are examples of transactions that may be immediately tracked by a TPS. Often, technology like barcodes and scanners makes tracking such transactions quicker, cheaper, and more detailed than it would otherwise be.
- Some other common terms we will define in more detail later in the course:
    - MIS – "Management Information System"
    - DSS – "Decision Support Systems"
    - ES – "Expert Systems"
    - EIS – "Executive Information Systems"
- An *Interorganizational System* (IOS) connects two organizations – for example, it may allow a company to automatically share inventory and backlog data with suppliers or customers.
- *Electronic Commerce* or *E-Commerce* refers to sales transactions in which at least one side of the transaction (buyer or seller), and perhaps both, is performed by a CBIS without direct human help.

Class 2 – Role of the Information Systems Department (ISD)

- Modern computer and network hardware software requires specialized skills and knowledge, at least for firms beyond a certain size.
- This means that the organization needs a sub-organization responsible for IT support: the Information Systems Department (ISD). Names vary from organization to organization; for example, at Rutgers it's called RUCS.
- In the early days of CBIS, ISD's "owned" *all* the information infrastructure because nobody else could understand it.
- As computers became more pervasive and user-friendly, managing IT resources has become a cooperative venture between ISD's, the departments/functional areas, and "end users" (individuals). Drawing the lines of who is responsible for what can be tricky.
    - At Rutgers, for example, RUCS operates the central network infrastructure, certain key systems like Eden and the outgoing mail server, and works in a consulting/contracting role to support other sub-organizations.
    - Individual or departmental management of resources will tend to be more responsive and understand user/departmental needs better
    - Central management will tend to have a better understanding of the technology in general, may promote better integration between departments, and can lead to economies of scale

    - This is a generic management issue that applies to lots of areas besides IT. For example, should each product division have its own product development engineering department, or should multiple divisions share an engineering department?
- In many organizations, the ISD has evolved (or should evolve) into a "business partner" with other departments, and not just a support organization – see below.

Class 2 – Competing in the Digital Economy (RTP Section 1.3 and Chapter 2)

The "digital economy" described in RTP Chapter 2 has certainly benefited from advances in computing power and information storage. But in the last 10-15 years, the most critical ingredient has been *networking* – the interconnection of multiple computers, and specifically the internet.

A brief history of internet technology (we'll do a bit more later in the course):

- The key to the present internet is the "TCP/IP" network technology developed in the 1970's
- TCP/IP was in wide use by research/academic US computer users by the mid 1980's. Typical applications were remote computer use (TELNET), file transfer (FTP), and e-mail
- In the late 1980's, significant research/academic use of TCP/IP began outside the US
- In the early 1990's, some physicists and computer programmers developed a network-browsable "hypertext" interface called the "world wide web"
- By the mid 1990's, the WWW drove a massive explosion in internet connectivity and applications; e-mail use "came along for the ride"
- The basic WWW interface was enhanced to gather as well as distribute data
- Technology was developed to link websites to databases
- This technology allowed sales transactions to occur over the WWW
- Physical products are typically delivered by package delivery networks like FedEx and UPS, which experienced symbiotic growth with the internet during the 1990's
    - The idea of a high-performance package delivery network using aircraft was itself pioneered by FedEx in the early 1980's.
    - IT tools have also been critical to the growth of package delivery networks
- For "digitizable" products like software and music recordings, the product itself could also be *delivered* over the network.

Computer/network-mediated business transactions are called *e-commerce*. Note that the e-commerce is very young; it barely existed a dozen years ago.

- Note that e-commerce does not *require* the WWW, although the WWW is a common foundation used to support e-commerce.
- For example, a "B2B" (business to business) e-commerce application in which one firm's information system automatically communicates parts orders to a supplier's information system would probably not use the WWW.

- On the other hand, "B2C" (business to consumer) e-commerce applications almost always use WWW interfaces.

Some common terms:

- *The internet*: the global network environment. Literally, it means a collection of interconnected networks.
- *Intranet*: the network within an organization; typically, it refers to portions of the network not accessible to those outside the organization.
- *Extranet*: one or more interconnected intranets, bridging multiple organization, but not openly accessible to those outside. For example, a firm might form an extranet with its dealers or key suppliers, in order to share critical inventory or product lead time information. This information would not be accessible or even detectable just by "googling".

Some examples of IT issues offering critical opportunities or challenges to companies:

- Obvious examples such as Amazon.com (my own observations; not in this part of RTP)
    - In some cases, a firm's entire business model is based on IT and the WWW
    - Consider Amazon.com, in their original business of selling books. The key observation is that there are a huge number of different books published, and many books appeal to a "thin", widely dispersed audience. Thus, any physical store must either limit its selection to better-selling books or carry a huge, slow-moving inventory.
    - The WWW provides an efficient way for customers to browse a huge selection (without mailing out gigantic catalogs)
    - Inventory can be concentrated in relatively few locations, where it turns over relatively quickly
    - Information systems streamline the "picking" and shipping of orders
    - Delivery via efficient package network carriers
    - Amazon.com simply could not exist without modern IT. Another example of such a business is Google.
- *Bringing 7-Eleven out of bankruptcy* (RTP pp.13-14)
    - Old supply system was chaotic:
        - Each store could have more than 80 deliveries per week, each with different items
        - Deliveries could occur during peak shopping hours and disrupt sales.
    - New supply system with stronger IT component:
        - Handheld computers used to place orders
        - Distribution centers consolidate each store's orders into a single 5 AM delivery the following morning
        - Real time sales and ordering data available to store managers and their superiors
        - Note that this application also involved redesigning the firm's *supply chain*. It's possible to improve supply chains without upgrading information systems, but IT can help a lot.

- *Sarbanes-Oxley at Blue Rhino* (RTP pp. 15-16)
    - Leading supplier of propane canisters for gas grills etc., sold and collected by independent local distributors.
    - Since market capitalization exceeds $75 million, the recent Sarbanes-Oxley Act requires both CEO and auditors to certify the firm's financial system and its controls
    - Accounting staff had to plug receivables and payables information from distributors into spreadsheets manually, on a monthly basis
    - Resulted in at least one week per month when inventories were not tracked accurately, so the firm had to carry an extra inventory "cushion"
        - Classic story at companies that have outgrown desktop computer tools.
        - Personal productivity tools like Excel and Word are great, but they are easy to outgrow
        - If many people are repetitively using the same spreadsheet or document, or it is used for routine, cyclic tasks like logging monthly or weekly sales, you have outgrown your desktop tools, and…
        - You should invest in a larger scale IT solution constructed with database technology
    - HR had to manually communicate information about new hires to the IT department (the process required manual intervention at both departments)
    - Purchasing required filling out manual forms
    - Sarbanes-Oxley accounting controls required improvements to all these systems; more efficient operations were a beneficial side effect.

(Note: rest of this class had to be abbreviated for lack of time; please refer to the cases mentioned below for some more illustrative examples.)

- A very similar situation is described in *Crown Media Complies with Sarbanes-Oxley* on RTP p. 51. Again, spreadsheet-based procedures were replaced with an accounting information system using database technology.
- Kmart and Sears: *Ignore Information Systems at Your Peril* (RTP p. 37)
    - While introducing ERP systems may be nightmarish, cobbled-together groups of loosely communicating legacy systems also have serious downsides
    - Information systems should not just be an afterthought in corporate mergers and acquisitions
    - Example: Kmart and Sears merged in 2004
    - Kmart had
        - 3 inventory management systems
        - 5 logistics management systems
        - 5 supply chain management systems (the difference between "logistics" and "supply chain" is not entirely clear here)
        - 4 purchasing systems
    - Sears had
        - 5 inventory management systems
        - 4 logistics applications
        - 5 supply chain systems
        - 6 merchandise planning systems

    - Each firm's IT infrastructure was individually a mess; now they have 37 systems to integrate.
    - IT considerations appear to be a significant obstacle to extracting value from the merger
    - While Sears/Kmart struggle with integrating these systems, WalMart and Target can press their advantage through more efficient operations and can get further ahead by adding new technology
    - Moral: the right IT tools can be a key "competitive advantage"
    - Consequence of moral: in many organizations, the IT department (ISD) should evolve away from just having a supporting role. Even in a business whose products/services are not directly IT-related, the IT department may need to evolve into a "partner" with (for example) marketing, finance, and/or operations.

Is the moral that all companies should invest aggressively in IT? Not necessarily:

- *Is Dollar General Really Thriving with Minimal IT?* (RTP p. 43): Dollar General has modern IT above the store level, but has minimized its investment in IT at the store level. Individual stores have only cash registers that capture transactions and upload them to the firm's central IT infrastructure once per night. The firm has had excellent financial performance, although "shrinkage" (theft and lost goods) are bothersome and hard to address due to the lack of in-store information systems.
    - *Possible moral*: if some aspects of IT are not critical to your business case, they may not merit aggressive investment.
- There are plenty of ERP horror stories to offset the success stories
    - 1999: Hershey reported a $19 million quarterly earnings drop when they brought ERP on line
    - 1999: Whirlpool was unable to ship large numbers of appliances after installing ERP
    - 1996/2001: FoxMeyer (a prescription drug distributor) blames bankruptcy filing on ERP introduction
    - ERP introduction often causes years of dislocation
    - Some firms simply "back out" of ERP introductions after spending millions of dollars.
- There is also no shortage of other IS efforts that have failed or have cost millions more than anticipated
    - 2002: installation of Security Audit and Analysis System (SAAS) at the IRS (RTP pp. 42-43)
    - The FAA introduced the National Airspace System Plan (NASP) to upgrade air traffic control information systems in 1982, but had made little progress by 1991 after huge investments. In 1991, the FAA introduces a more incremental plan that has been much more effective. For example, one element of the incremental plan involved (temporarily) programming new hardware to simulate old hardware which was becoming too difficult and expensive to maintain.
- IT missteps, especially grandiose ones, can be very expensive.
- Sometimes very ambitious, far-reaching IT upgrades that sound great are too difficult to implement or introduce in practice. Gradual upgrade and consolidation of systems may sometimes work better.

One clear message here is that it’s important to think about IT early and try to get it right.

- Changing and upgrading systems can be really painful
- But waiting too long to put in a system is also painful; countless firms are wasting countless employee hours fiddling with spreadsheets when they should have moved to a comprehensive, multi-user database solution long ago.

There’s no magic formula for how firms should approach IT, or whether a particular IT project makes sense. Near the end of the course, however, we’ll discus the “SDLC” methodology for evaluating, acquiring, and developing information systems.

I believe that basic technical understanding of the technology is key to making good decisions about it. That is why we will start working hands-on with relational database software in the next class.

Class 3 – Lab: Introducing MS Access tables and forms

Class 4 – Lab: Introducing MS Access queries and reports

Class 5 – Memory storage calculations

See *Memory Storage Calculations* handout.

Class 6 – Data Management (Excerpts from RTP Chapter 4)

We have seen a little bit now about the tables in which business data are stored, and to how to calculate the amount of storage they might consume.

Chapter 4 of RTP addresses issues of data management. Such issues include:

- How many tables should be in the database an information systems using?
- What data should be in each table?
- How the tables should be connected to one another?
- When an organization has more than one information system (and most do), what information should be in each system?
- How should the systems communicate?

The issues concerning interrelated and overlapping systems resemble in some ways the same questions with tables, but on a larger scale. In this course,

- We will get into a lot of technical detail about how tables should be organized within one database system. It’s a well understood area and the basic concepts are not too difficult.
- As the way entire “systems” should interact, we will take a more descriptive, superficial approach.

Case illustrating some common data management issues: RTP pages 98-99 – MetLife around 1999.

- The firm had a very large number of information systems, one for each product line

- Example: if a client informed the firm of a change of address when they updated their automobile policy, that change would not propagate to their life insurance policy
    - This is an example of how *redundancy* – storing the same data in more than one place – can be a headache for information systems (relate story of safety deposit box)
- Data from all these different systems was hard to integrate, making the firm hard to manage. This problem may have contributed to
    - Poor financial performance relative to competitors
    - Improper sales practices in some units, resulting in a scandal and almost $2 billion in fines and penalties.
- The firm wanted to better integrate business units in response to legislative changes removing barriers between the banking, insurance, and securities industries.
- The company reorganized the data systems serving individual customers with products like life, home, and auto insurance. All these systems now use a common customer database.

MetLife's problems were not unusual, and many companies have similar problems today. Generically, some common data management problems facing today's organizations are:

- The volume of data increases as new data is added. Historical data is often kept for a long time, so typically data comes in faster than it is deleted. New technologies mean that gathering new data is easier and faster. So, not only is the total volume of data increasing, but the rate at which it is increasing is also increasing!
    - Should a firm take advantage of every possible opportunity to gather data?
    - For example, any website can gather "clickstream" data: records of exactly how users moved around a website, whether they bought anything or not. This data may be of value, but can pile up quickly.
- Data tend to be scattered throughout the organization. Older organizations tend to have many "legacy" systems that communicated poorly, causing severe problems – the MetLife case is an example. Thus, it is often desirable to centralize data storage, but by no means always – it may be better to leave departments or working groups "in charge" of the data they use the most. It is costly and risky to replace older "legacy" information subsystems that are working smoothly. Replacing a large number of smaller systems with one larger one can often be very complicated and costly. Sometimes it may be better to created "federated" systems that combine information from constituent systems.
- Data accuracy – many organizations have far more errors in their databases than they are aware of. One cause is unnecessary data redundancy (see RTP page 101), but there are other causes too.
- We may also want to use data from outside the organization (either public-domain or purchased).
- It may also be advantageous to share some information with suppliers or vendors (for example, sharing information about inventories can reduce inventory fluctuations and costs throughout a "supply chain").
- Data security and quality are important, but are more easily jeopardized the larger and more complicated an information system becomes.

We can classify the processing tasks information systems perform as follows:

- *Transactional processing* (sometimes called TPS): keeping track of day-to-day events, such as logging orders and shipments, and posting entries to accounting ledgers. In terms of a data table, transaction processing means an ongoing process of adding rows (for example, to reflect a new order), modifying table cells here and there (for example, if a customer changes their telephone number), and perhaps deleting rows.
- *Analytical processing*: means combining data from many table rows in order to obtain "higher-level" information. Entering a row into a table to reflect a newly-received order would be transaction processing; an example of analytical processing would be computing the number of orders and total dollar value of orders for this month, and comparing them to last month.
    - Analytical processing can be as simple as sorting, grouping, and summary calculations in an Access query or report. For example, providing all group managers with a summary of their groups' costs for the month, broken down by cost category. This kind of application can be called "classic" MIS (Management Information Systems).
    - Analytical processing can get a lot more sophisticated. For example, *data mining* refers to using sophisticated statistical or related techniques to discover patterns that might not be apparent in standard reports.
    - *Decision support* can involve using the data to help managers make complex decisions, *i.e.* how to route 400 shipments from 20 warehouses to 100 customers.
    - Database systems like Access don't ordinarily do data mining or decision support by themselves. For such uses, they usually need to be connected to other pieces of software.

Sometimes it can be mistake to do transaction processing and analytical processing on the same database, especially if the analytical processing is very time consuming or complex.

- The analytical processing may make the transaction system run slowly
- Conversely, the transactions may interfere with the analytical processing and make *it* run to slowly
- If an analytical processing step takes too long, the data it is using may change in the middle of its calculation. "Locking" the data to avoid this can block transaction processing.

It may be better to make a copy or "snapshot" of the database used for the transaction system.

- This is often called a "data warehouse" – see RTP Section 4.4.
- You can do a lot of analysis on the data warehouse without disrupting the transaction system and a lot of transactions without disrupting data analysis
- The data warehouse will not reflect the very latest transactions, but for large-scale aggregate analysis, that may not be a big problem.
- Another reason to create a data warehouse might be to consolidate information from several different transaction systems so it can be analyzed together (see Figure 4.9, RTP page 112).

At this point, we will start getting into the details of data management for the case of a single system. You may think of this material as a very detailed expansion of the subject matter on RTP Sections 4.2 and 4.3.

Remainder of Class 6 – Roughly follows GB pages 218-222

(However, I also introduced the notion of a *repeating group*.)

Classes 7-11 – Database design; multiple tables in Access

Refer to class handouts.

Classes 12-13 – Network computing

Note: these classes cover roughly the material in RTP chapter 5, including its appendix. However, my emphasis will be different, and I'll give some details not in the text.

Broadly speaking, in terms of major impact and commercialization:

- The 1980's was the decade of personal computing (single-user, desktop computers)
- The 1990's was the decade of networking

Many of the relatively recent advances and technologies we now associate with computers are really networking-based.

Basic definitions:

- Something that connects two computers is a *link*
- Many computers connected by many links comprise a *network*.
- Each computer on the network is called a *node*.
- Generally speaking, data should be able to get from any network node to any other node.
- There are many different shapes (or topologies) that can be used to build a network
    - Today the predominant network topology is a collection of interconnected "stars" or "buses"
        - "Star" – each computer has an individual connection to a central interconnection point
        - "Bus" – a group of computers share a common communication channel (either a radio frequency of a metal wire). Only one can send information into the channel at any given time.
    - At one time, interconnected "rings" were also popular, and they are still in use (see Figure 5.1).
- Some nodes on the network are specialized computers that serve primarily as connection points whose job is to make sure data gets sent to right place. Some common names:
    - Switches
    - Hubs
    - Routers

  Most of this equipment is not particularly visible to most people, although a lot of people now own routers or wireless routers for their homes.

Kinds of links

- Link speed is usually measure in *bits* per second (b/s), with the usual prefixes K (kilo/1,000), M (mega/1,000,000), G (giga/1,000,000,000), etc. Note that network engineers usually use the *decimal* form of these prefixes, not the binary ones.
- Wires (usually copper) these can be used in many ways.

    - Twisted-pair wires, such as
        - Regular telephone wiring
        - Ethernet wires, which currently come in three flavors, 10 Mb/s, 100 Mb/s, and 1Gb/s.
    - Coaxial cables, like those used for cable TV. These have higher theoretical capacity but are harder to work with (they are stiff and hard to connect together).

  For wires, there is a trade-off between distance and the number of bits per second that may be transmitted. You can transmit with very high data rates, or over very long distances, but not both. This trade-off arises from the basic electrical resistance properties wires.
- Optical fiber (carries light pulses)
    - Invented about 30 years ago
    - Much better than wire for combining high data rates and long distances. Links can have capacities in the many Tb/s
    - More difficult to work with than either twisted-pair wires or coaxial cables. In particular, it's hard to "splice" two of them together (however, that also makes them more secure).
- Unconfined electromagnetic waves (radio/infrared/microwave) – "wireless"
    - Microwave links (can be directional – can be formed into a narrow beam)
    - Satellite links
    - Within-building ("wi-fi") broadcast: current capacities typically 1-54 Mb/s
    - Wider area wireless: slower than wi-fi right now (cell-phone modems); this technology is just emerging, but I believe that demand will make it successful. Current services are advertised at $80/month for 400-700 Kb/s access, with approximately the same coverage area as standard cell phones.

A quick history of computer communications:

- The first large-scale electronic networks built were telephone networks. But they were not used by computers initially, because computers didn't exist. In fact, "computer" was actually a job title for a person who did mathematical calculations for engineers and scientists.
- When businesses started using computers, each organization had its own computer in its own room. Data got in and out of the computer room by being physically carried as punched cards, printed reports, magnetic tape etc. (eventually floppy disks, too) – later called "sneakernet".
- People began placing I/O devices outside the computer room, connected by wires: printers, card readers, terminals (=printer + keyboard or screen + keyboard), etc.
- Technology was invented to encode (modulate) data into sounds the telephone network could carry. The data would be "demodulated" back into bits at the other end (thus the term "modem" – *mod*ulator/*dem*odulator); see TRP p. 461.
    - This allowed people to have terminals at home and work over telephone lines
    - Many other business applications involving sending or receiving data from remote locations
    - Early modems were slow (100 b/s = 0.1 Kb/s in the 1960's). This gradually increased to about 56 Kb/s today.
    - The technology is still widely used, but in decline

- In the late 1960's, interest was growing in large general-purpose data networks independent of the telephone network.
    - Before, these existed only for specialized application (mostly military)
    - ARPANET – the (defense department) Advanced Research Projects Agency NETwork was built in the early 70's
    - This network gradually evolved into "the internet"
    - The internet had a fairly small user base until the mid 80's. Then it began to gather momentum
    - In the early 90's, the "world wide web" application of internet technology became very popular and drove a massive expansion of the internet (along with the ".com" boom)
    - In the 90's there was a general telecommunications boom of which the internet boom was a big part. Numerous firms tried to secure their place in the boom by building lots of network links, especially in North America
    - A lot of network capacity was built. About the same time technology appeared that greatly increased the amount of data an optical fiber could carry (by simultaneously sending multiple information streams using light beams of different colors). The result was overcapacity and a major telecommunications business "crash". However, internet use continues to climb.

How networks work: LAYERING is very important

- Bottom: physical layer – the physical workings of the links (wire, fiber, wireless, etc.)
- Network layer (typically "IP", standing for "internet protocol"): handles the identification of particular computers on the network, and describes the structure of the network.
    - Currently, each computer has a 32 bit "IP address" (usually split into four bytes printed in decimal like 128.6.59.202).
    - The addresses have structure – for example "128.6" in the first two bytes of the address means somewhere at Rutgers (although 165.230 also means Rutgers), the 59 designates a particular "subnet" (roughly the same as a building or part of a building), and the 202 identifies which computer on the subnet.
    - Note that most computers also have a "hostname" that is easier for humans to remember, like "business.rutgers.edu" or "www.amazon.com". While these hostnames are related to IP addresses, they aren't exactly the same. Special computers on the network, called "name servers" provide the translation. Small organizations may not have a name server, relying on a name server from their internet service provider. Large organizations like Rutgers may have dozens of name servers.
    - The end of the hostname is often called the "domain name". For example, the domain name in "business.rutgers.edu" is "rutgers.edu". The domain name usually indicates which institution owns the computer, and what kind of institution it is. "Rutgers.edu", for example, means the computer is at Rutgers, which is an educational institution.
    - 32 bits are no longer enough space for an IP address, and we will gradually move from IPv4 (32 bits) to IPv6 (128 bit addresses). Various workarounds suffice for now:

        - Dynamically allocating IP addresses only when computers are actively connected to the network ("DHCP" is a common way of doing this), or
        - Grouping small sets of computers to share a single IP (network address translation or "NAT"). For example, if you have a router in your house or dorm room, it is probably doing NAT translation, so that all the computers in your house effectively share a single IP address.
- Transport layer (typically "TCP"). Specifies the structure of the information being moved across the network
    - TCP specifies up to 64K (binary K) logical "ports" for each computer on the network. Each port is typically used for a different application. For example, standard web page viewing usually uses port 80.
    - For each port, there may be one or more "sessions" or logical connections between to computers. For example, you could have two independent web browser windows connected to the same website from your own PC; each of these windows would represent a different session for TCP port 80 on the server. Each session is initiated by one computer sending a request to another. If the second computer agrees, the session is opened, and data may then flow in either direction.
    - For each session, there may be a sequence of messages in each direction
    - TCP is a "packet switched" protocol – messages are cut up into "packets" that might take different paths through the network and are reassembled at the destination. By contrast, telephone networks are "circuit switched" – the whole conversation uses the same route through the network. The size of packets varies by the application and network hardware in use, but they are typically roughly on the order of magnitude of 1KB.
- Application layer: specifies different protocols to move data for various uses. These protocols constitute an "alphabet soup":
    - First: TELNET (old) – run a terminal session (a back-and-forth text-based interaction between a person and a computer – kind of like the command prompt in Windows)
    - FTP (old) – move files back and forth (still in some use when security isn't an issue)
    - SSH – encrypted terminal sessions and file transfers. This is how you connect to the "Eden" system to do text-based interactions. This accomplishes the same basic tasks as TELNET and FTP, but is far more secure. Other protocols can be "tunneled" through SSH to make them secure.
    - HTTP/HTTPS – hypertext transmission. This application appeared in the early 1990's and rapidly evolved into a way of projecting a wide range of graphic user interfaces across the internet. The "S" in HTTPS means secure/encrypted. HTTP is a much easier and more secure way to do arbitrary things on a remote user's screen than making the user run custom software.
    - SMB, NFS – file sharing. Makes disks on a distant computer look like they're on yours.
    - SMTP – sending e-mail to and between mail servers (computers that can route e-mail). This is a "push" protocol: the computer requesting the connection sends the messages; when your computer sends an e-mail, it typically uses SMTP.

    - POP3, IMAP – retrieving mail from e-mail servers. These are “pull” protocols: the computer requesting the connection receives the messages (if there are any)
    - And many, many, more…
- Typically, each protocol uses a single TCP port (or perhaps a few). For example, HTTP usually uses port 80, and SSH usually uses port 22.

Some more notes on layers and protocols

- The standard URL (“Uniform Resource Locator”) used by web browsers has the basic form

  PROTOCOL://*hostname*/*additional-information*

  Thus, http://www.rutgers.edu/ means “HTTP protocol, computer with hostname www.rutgers.edu, no additional information”. A numeric IP address can also be used instead of a hostname, as in http://123.57.12.92/obscure-stuff (but is a clue that the site is probably not be safe to connect to).
- As you move downwards in the protocol layer “stack”, more and more “bookkeeping” data — also called “headers” — get appended around the data you are sending. This means the actual number of bits transmitted can be substantially more than you might think. For example, when a message is divided into packets, information is added to each packet so that it gets routed correctly and the packets may be assembled again in the right order.
- TCP and IP are usually used together and are known as “TCP/IP”
- You can run more than one network layer on top of a physical layer on the same link (for example, IP and AppleTalk)
- You can run several transport layers on top of a network layer (for example, TCP and UDP on top of IP)
- And, of course, you can run many application layers on top of a transport layer (SSH and HTTP on top of TCP)

Kinds of networks

- LAN – “Local Area Network” on the level of a single building, part of a building, or office
- WAN – “Wide Area Network” a somewhat vague term for a network that covers a “non-local” geographic area, that is, something larger than a LAN. An example would be the network that connects Rutgers’ various campuses (including Camden and Newark).
- Enterprise network – refers to the totality of an organization’s networks, both its LANs and its WAN(s) together.
- Internet – multiple networks networked together
    - The idea of *an* internet preceded the current notion of *the* internet – “the” internet came into existence when almost everything got connected into one huge network!
    - The “IP” network layer was specifically designed to make it easy to create internets. That is why all internets essentially merged into “the” internet that grew so quickly in the 1980’s and 1990’s, and conversely why IP is now the dominant network layer.
- Intranet – the portion of an organization’s enterprise network that is not accessible by arbitrary internet users

- Extranet – when multiple organizations connect their networks in a way that is not fully accessible from outside that set of organizations
- VPN – "Virtual Private Network" – an intranet or extranet that physically uses the general internet, but is encrypted in such a way that it looks like a private WAN that outsiders can't snoop on (we hope).

Current network technology

- Most firms now have LANs implemented with copper wire, usually Ethernet, and now also building-level wireless
- Many larger firms have WANs containing wire and/or fiber and maybe some satellite or microwave links (depending on the firm's size). The longer links in these networks are typically leased from ISP's (see the next item)
- *Internet service providers* (ISP's) are firms maintaining interconnected, overlapping networks made primarily of fiber (examples: AOL, ATT, Sprint, etc.), but possibly also involving wire, satellite, and microwave links. ISP's also lease capacity to firms for use in WANs. Large- and medium-size firms connect directly to ISP's.
    - Also, there are some non-profit alternatives to ISP's, like "Internet2", a consortium of large universities like Rutgers, and other large nonprofit organizations.
- Large firms can afford to lease dedicated high-speed connections to ISP's, like "T3" lines
- The dreaded "last mile": smaller firms and individual households connect to the ISP's in various non-ideal ways:
    - By phone and modem (sometimes directly to an employer instead of ISP)
    - Cable modem – signals carried over the same coaxial cable that distributes TV signals. Capacity usually 0.5-7 MB/s, but capacity may be shared with other users in the neighborhood
    - DSL – signals carried over regular phone lines, but not at audible frequencies. About 0.5-1 Mb/s, but occasionally faster. Only works if you are within 2 miles of telephone switching center, but does not have capacity sharing problems.
    - Cable to the home (example: Verizon FiOS) – run optical fiber all the way to individual houses. This is the most straightforward technology, but is being rolled out slowly in limited areas. Depending on the price paid, individual households can get 5-20 Mb/s download speeds and 1-5 Mb/s upload.
    - WiMax – (see RTP pp. 223-224) a wireless networking technology with a theoretical reach of up to 70 miles and transfer rates up to 70 Mb/s. However, there is a tradeoff between distance and speed, and more typical performance would be about 10 Mb/s over 6-7 miles, if a "line of sight" connection is possible. This technology is rare in the US and is currently targeted at areas lacking wired infrastructure (especially rural areas). Receiver units can be quite expensive ($349 is mentioned on RTP p. 225).
- Most network connections carry a fixed charge per month, without tracking the exact number of bits sent – one reason we have so much "spam"!

Uses for networks are expanding all the time. For a fairly current catalog, see Section 5.4 of RTP.

Data transfer calculations – how much time will it take to move a data file? (The file could contain text, data tables, video, audio, etc.)

- Calculate the amount of data to be moved in *bits*
- Divide by the speed of the transmission line in bits per second
- Convert from seconds to larger time units like minutes or hours if necessary
- Remember:
    - File sizes are commonly in bytes, most often with *binary*-style K, M, G etc.
    - Transmission line speeds are usually in *bits* per second, with *decimal*-style K, M, G etc.
    - This mismatch is annoying, but is the common convention.
    - It's easiest to convert the file size to decimal bits, and then divide by the line speed.
    - The protocol stack will add header information that will cause the real download to take longer
    - Network congestion or malfunctions could cause even more delays

Sample file transfer calculation: Suppose we want to do "video-on-demand" downloads of 4 GB movies in DVD format (binary-style GB). How long would that take over a 1 Mb/s DSL line?

Size of movie = (4 GB)(1024^3 B/GB)(8 bits/B) = 3.44×10^{10} bits

Seconds to transfer with DSL = (3.44×10^{10} bits)/(1×10^6 bits/sec) = 3.44×10^4 sec
= (3.44×10^4 sec)/(60 sec/min × 60 min/hr) = 9.54 hours – probably not acceptable!

Note that actual transfer times would be somewhat larger due to overhead (headers) added by the application, transport, network, and physical network layers, and because of possible network congestion.