Tutorial at the AICCSA '2007)
May 13-16, 2007, Amman, Jordan

Control Problems in Multi-service, Multi-platform Telecommunication Networks

Instructor: Prof. Franco Davoli
Department of Communications, Computer and Systems Science (DIST),
and Italian National Consortium for Telecommunications (CNIT) – University of Genoa Research Unit
University of Genoa
Via Opera Pia 13
I-16145 Genova
Italy

Tel.: +390103532732
Fax: +390103532154
E-mail: franco@dist.unige.it

Summary: Telecommunication Networks are complex, distributed, large-scale systems. Networking is basically a resource allocation activity, in a broad sense, where resources are typically represented by bandwidth, storage and processing capacity. The optimization of these resources, whose main final goals are to provide Quality of Service (QoS) to the users at reasonable and competitive prices and to maximize the revenue of the network operator, can be viewed under two different, related perspectives: for network planning purposes and for real time control of network operations. The importance of these aspects is enhanced by the ever-increasing presence of transfer modes based on statistical multiplexing paradigms (e.g., the Internet) and of multiple services within all types of networks. Moreover, complexity is added to the problem by the heterogeneity of networking platforms: though a sort of common paradigm at the network layer and above is that of the Internet Protocol Suite (and related QoS mechanisms), there are a number of different physical transport environments, which have widely different characteristics in terms of transmission capacity, error resilience, operational complexity, and scalability. The aim of the tutorial is to explore different areas in networking (QoS-Internet, cellular networks and wireless LANs, satellite networks, optical networks) from the point of view of resource allocation and QoS control, to outline the main problem areas and to point to some common control techniques arising in the different environments.

Outline of Topics:

Control Problems in Telecommunication Networks

- General aspects and common ground
- Problem areas in networking control
- Heterogeneous networking environments
- Quality of Service
- Timescales - control, management, planning
- Networking technologies

Optimal Control of Dynamic Systems

- Representations of dynamic systems
- Controlled Markov chains
- Markov Decision Processes
- Functional and parametric optimization
- Dynamic Programming
- Optimization Techniques

Call Admission Control, Bandwidth Allocation, Congestion Control

- Admission policies
- Service separation - Decoupling low- and high-level constraints
- Bandwidth allocation
- Dynamic routing of flows
- Pricing
- Applications in wired and wireless networking platforms

Instructor's short bio: **Franco Davoli** received the 'laurea' degree in Electronic Engineering in 1975 from the University of Genoa, Italy. Since 1990 he has been Full Professor of Telecommunication Networks at the University of Genoa, where he is with the Department of Communications, Computer and Systems Science (DIST). From 1989 to 1991 and from 1993 to 1996 he was also teaching classes in Telecommunication Networks at the University of Parma, Italy. His past research activities have included adaptive and decentralized control, large scale systems, routing and multiple access in packet-switched communication networks, packet radio networks. His current research interests are in bandwidth allocation, admission control and routing in multiservice networks, wireless mobile and satellite networks and multimedia communications and services. He has co-authored over 250 scientific publications in international journals, book chapters and conferente proceedings. In 2004, he has been the recipient of an Erskine Fellowship from the University of Canterbury, Christchurch, New Zealand, as Visiting Professor. He has been Principal Investigator in a large number of research projects, and has served in several positions in the Italian National Consortium for Telecommunications (CNIT), including the direction of the National Laboratory for Multimedia Communications in Naples in the period 2002-2004; he is currently Vice-President of the CNIT Management Committee. He is a Senior Member of the IEEE.

Target Audience: students, researchers, instructors and professionals interested in the field of telecommunication networks. A minimum knowledge of basic networking principles is required.

cnit Italian National Consortium for Telecommunications

Control Problems in Multi-service, Multi-platform Telecommunication Networks

SAMPLE SLIDES

Franco Davoli
Department of Communications, Computer and Systems Science (DIST), University of Genoa, Italy

CNIT-University of Genoa Research Unit

franco@dist.unige.it

Outline

- *Control problems in Telecommunication Networks*
  - General aspects and common ground
  - Problem areas in networking control
  - Heterogeneous networking environments
  - Timescales - control, management, planning

Outline (cont'd)

- *Call Admission Control, Bandwidth Allocation and Routing, Congestion Control*
  - Admission policies
  - Service separation - Decoupling low- and high-level constraints
  - Bandwidth allocation in ATM and IP terrestrial, mobile wireless and satellite networks - Dynamic routing of flows
  - Scheduling
  - Pricing

Outline (cont'd)

- *Instances of control techniques in specific environments*
  - Access networks
  - Mobile wireless networks and satellite networks
  - Cross-layer approaches
  - Approximation techniques and parametric optimization

Introduction

- A large number of dynamic control and resource allocation problems arise in almost all types of communication networks.
- Among others, some examples of the most commonly found ones are:
  - Connection Admission Control (CAC)
  - Bandwidth Allocation
  - Congestion Control
  - Routing
  - Scheduling
  - Power control in wireless networks

Introduction (cont'd)

- Such problems are encountered in different networking environments, cabled and wireless; basically
  - TDM-based structures (circuit-switched telephone networks; mobile radio networks, even in conjunction with CDMA; satellite networks)
  - ATM networks
  - IP networks with DiffServ/IntServ paradigms, MPLS
  - Optical networks (with MPλS, GMPLS)
  - Wireless networks
- The various structures may appear together (in particular, IP-over-X)

Introduction (cont'd)

- Control problems appear at various architectural layers
  - Physical and Data Link: dynamic power control and fade countermeasures, bandwidth allocation among users and services, multiple access, …
  - Network: dynamic bandwidth allocation, routing, Call Admission Control, packet scheduling, …
  - Transport: elastic bandwidth allocation, congestion control, …
  - Application: congestion control (e.g., TCP-friendly applications), rate adaptation, pricing, …
- *Cross-layer approaches* (i.e., exploiting information from other layers for control purposes) are often advisable, especially at lower layers in noisy environments.

Purpose of the Tutorial

- To give an overview of some control issues and techniques commonly used in telecommunication networks
- To show instances of their application in different networking environments
- To point out possible open problems and research areas

Goals of control

- The ultimate goal is to provide some level of *Quality of Service* (*QoS*) to the entities (data units, connections, applications, end users) that are being considered, depending on the specific layer or the "granularity" (or on the scope or "width") they are looked upon.
- According to ITU-T E800, QoS is interpreted as
  - ✓ The overall effect of performance-enabling services that determine the degree of satisfaction of a service user.
  - ✓ From the viewpoint of the telecommunication network, QoS translates into the capability of the network to guarantee a specific service level.

Quality of Service

- Indeed, the term QoS has a number of interpretations, which range from the quality perceived by the service user to a set of performance (in general, layer-specific) parameters that is necessary to specify to obtain the desired level of service.
- E.g., one may distinguish
  - **Intrinsic QoS** : directly provided by the network and described in terms of objective indicators, like loss (of data units or connections) and transfer delay.
  - **Perceived QoS - P-QoS**: as subjectively measured by the *Mean Opinion Score (MOS).*
  - **Assessed QoS** : as referred to the user's willingness to continue using a service. Related to P-QoS, but also dependent from the pricing mechanism, the support guaranteed by the provider and other commercial and market aspects.

Service Level Specification

- QoS provisioning is often offered in terms of objective indicators, by using a *Service Level Specification - SLS*.
- The SLS is a set of performance indexes and of their required values that together define the service offered to a given traffic.
- The SLS is the technical part of an agreement, negotiated between service user and provider, relatively to the characteristics of the service itself and to the associated set of metrics (*Service Level Agreement - SLA*).

Applications

- Which applications need some form of QoS?
- All applications requiring a specified level of "guarantee" from the network
  - Services for the transport of aggregate information (bandwidth from providers, VPN)
    - In the access network
    - In the backbone network
  - Videoconferencing, videotelephony
  - VoIP, Internet Telephony
  - Tele-medicine
  - Tele-education
  - Remote Control
  - Emergency (Disaster Recovery) applications
  - ...

Requests for QoS

- Market requests
- Widespread diffusion of the *Internet Protocol Suite* as a “universal” platform
- Need of mechanisms to provide quality *end-to-end* QoS on IP networks and across multiple domains.

ITU-T QoS Classes (for IP)

ITU-T Y-1541

QoS Class	Characteristics
0	Real-time, delay jitter sensitive, highly interactive
1	Real-time, delay jitter sensitive, interactive
2	Data transactions, highly interactive
3	Data transactions, interactive
4	Low loss (short transactions, streaming data flow)
5	Traditional applications of *best-effort* IP networks

QoS metrics (IP)

- *IPLR - IP Packet Loss Rate*
- *IPTD - IP Packet Transfer Delay*
- *IPDV - IP Packet Delay Variation*
- *IPER – IP Packet Error Rate*
- *Skew* (average value of the delay difference among packets belonging to different, mutually synchronized, media)
- ...

cnit

QoS metrics - Requirements

		QoS Classes					
Performance Parameter		0	1	2	3	4	Class 5 *Un-specified*
IPTD	Upper limit on average IPTD	100 ms	400 ms	100 ms	400 ms	1 s	*Un-specified*
IPDV	Upper limit on $1\text{-}10^{-3}$ quantile of IPTD less min IPTD	50 ms	50 ms	*Un-specified*	*Un-specified*	*Un-specified*	*Un-specified*
IPLR	Upper limit on packet loss rate	1×10^{-3}	1×10^{-3}	1×10^{-3}	1×10^{-3}	1×10^{-3}	*Un-specified*
IPER	Upper limit	1×10^{-4}					

cnit

QoS metrics - An example

MULTIMEDIA QoS REQUIREMENTS

Traffic Type	Maximum Packet Loss	Maximum One-Way Latency	Maximum Jitter	Guaranteed Priority Bandwidth per Session
Voice over IP	1 percent	200 ms	30 ms	12 to 106 kbps*
Videoconferencing	1 percent	200 ms	30 ms	Size of the session plus 20 percent
Streaming Video	2 percent	5 seconds	Not applicable	Depends on encoding format and video stream rate
Data	Variable	Variable	Variable	Variable

* Depending on sampling rate, codec, and Layer 2 overhead

Functionalities and Tools

- Identification of traffic flows
- Call Admission Control
- Traffic Engineering
- Scheduling (Service discipline)
- Flow and congestion control
- QoS Routing
- Resource Allocation

cnit

QoS control - Time scales

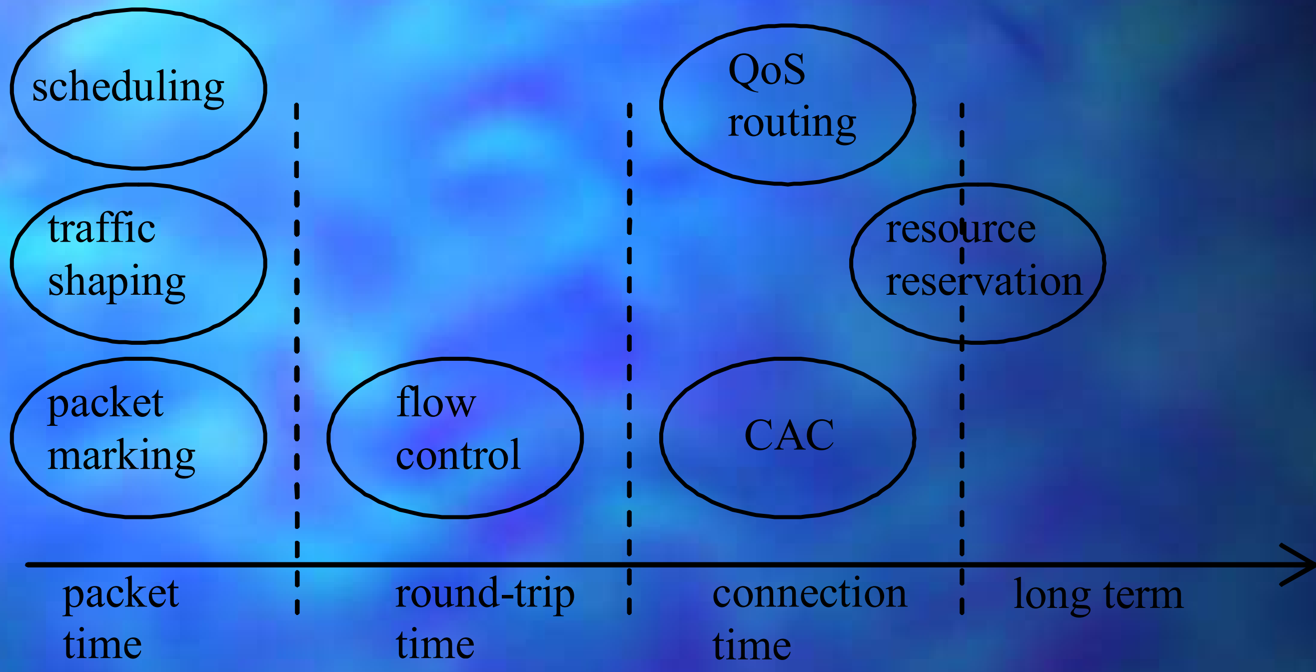

A wide range of time scales, with orders of magnitude from few μs to minutes, hours and days. Accordingly, a set of (related) resource allocation and control problems, spanning

- Network Control
- Network Management
- Network Planning

QoS mapping

- Services offered by lower layers should provide QoS *mapping* functions for the benefit of higher layers
- Implementing end-to-end guarantees (if at all possible!) would imply cooperation among layers
- However, care should be taken in cross-layer approaches, in order not to disrupt architectural principles that ensure interoperabilty

Transport technologies

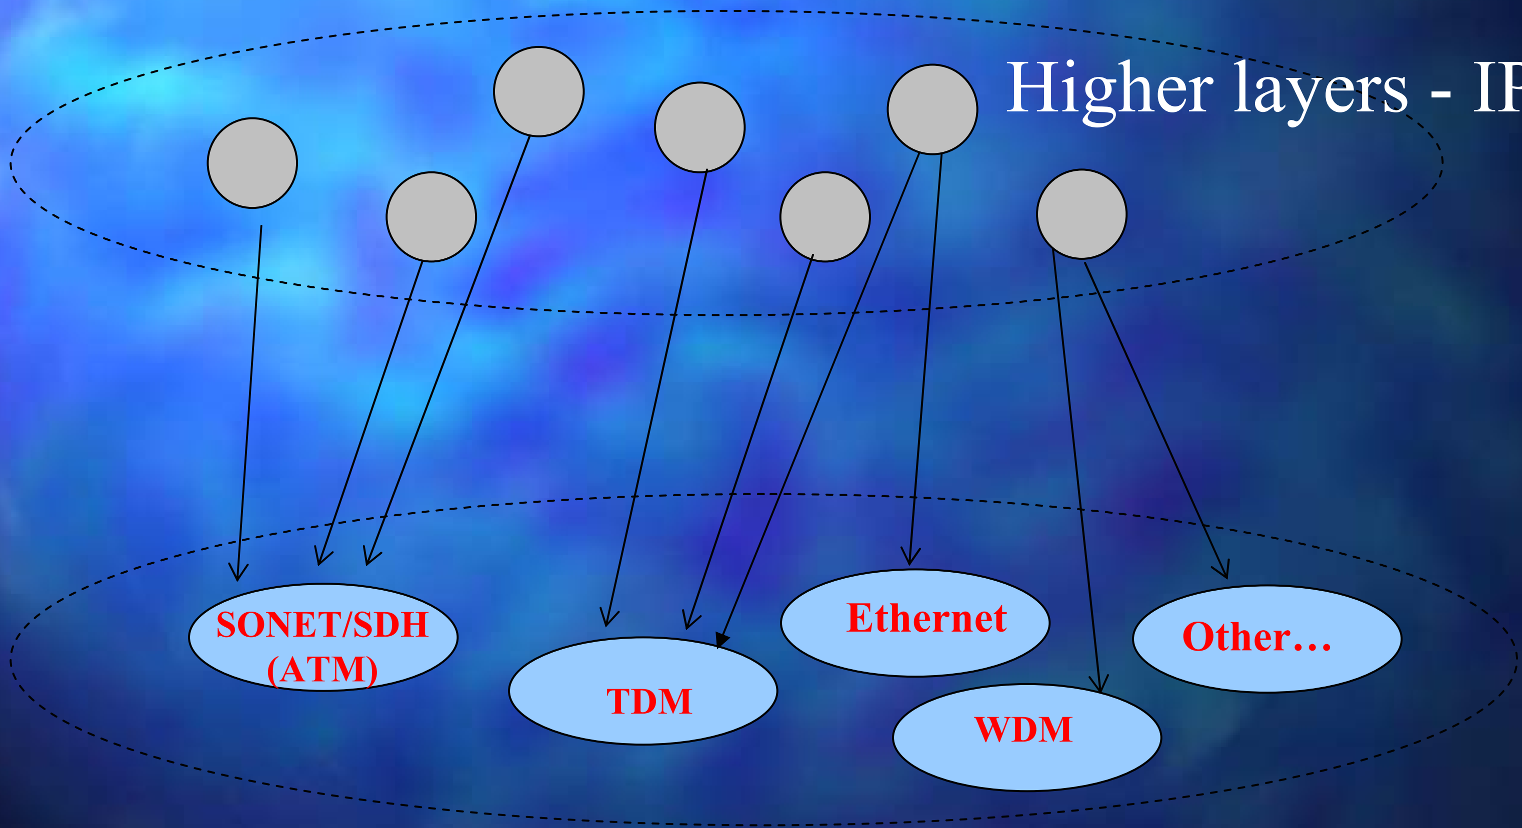

QoS control - Technologies

- "*Throwing bandwidth at the problem*"
- ATM
- IP (IPv4 and IPv6)
  - DiffServ / IntServ
- MPLS

Technologies - MPLS

IP

Simple
Universally widespread

MPLS

ATM

QoS handling
Control functions

Label Switching - Traffic engineering capabilities

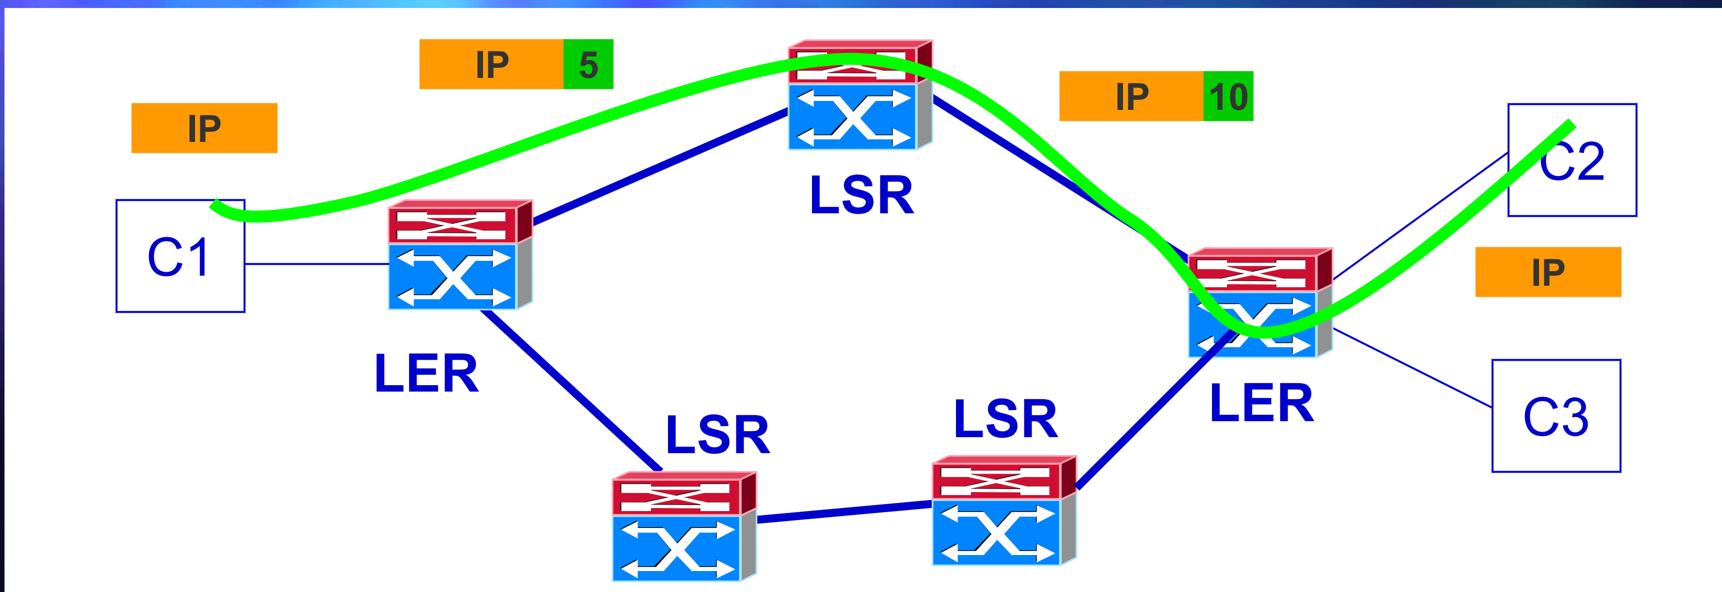

Technologies - MPLS

- *Label*
- *Experimental*
- *Stacking bit* (indicates presence of more *labels*)
- *Time to live*
  - IP on guaranteed performance network

cnit

Technologies - MPLS

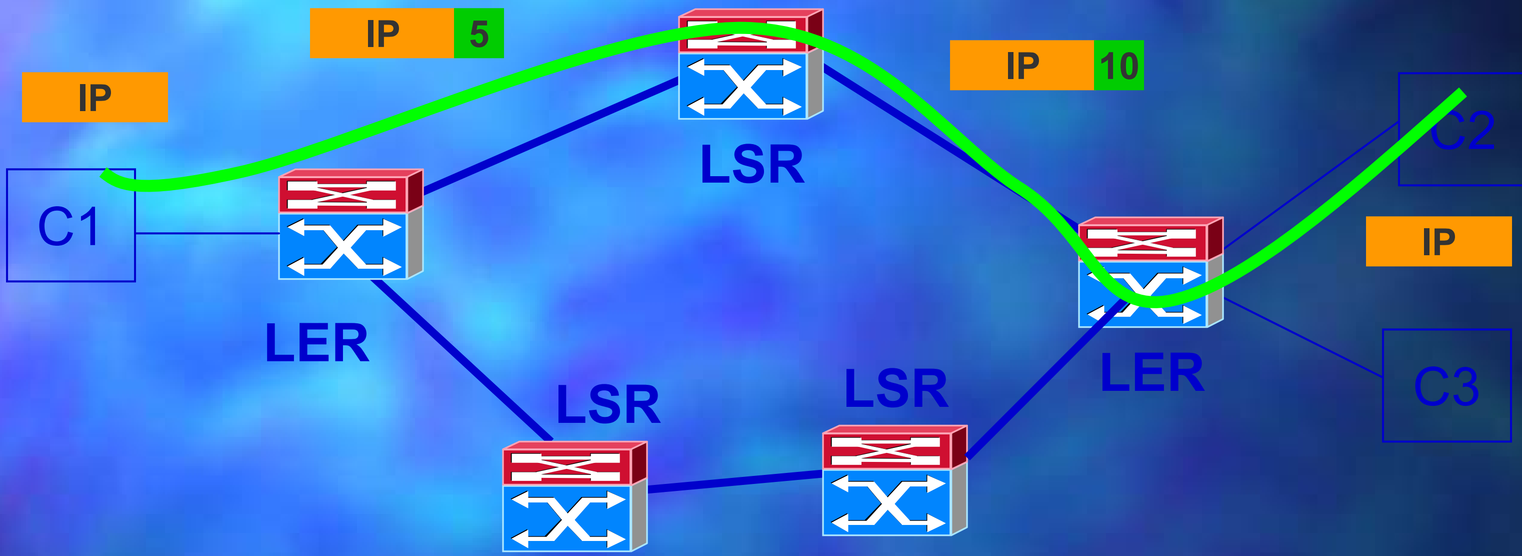

- LER (*Label Edge Router*) at ingress applies *Label* to packet and sends over correct LSP (*Label Switched Path*).
- LSRs (*Label Switched Router*) switch packet, swapping labels
- Egress LER eliminates label and forwards packet with IP forwarding procedure

Technologies - Flow ID

- ATM
- IPv4
- IPv6

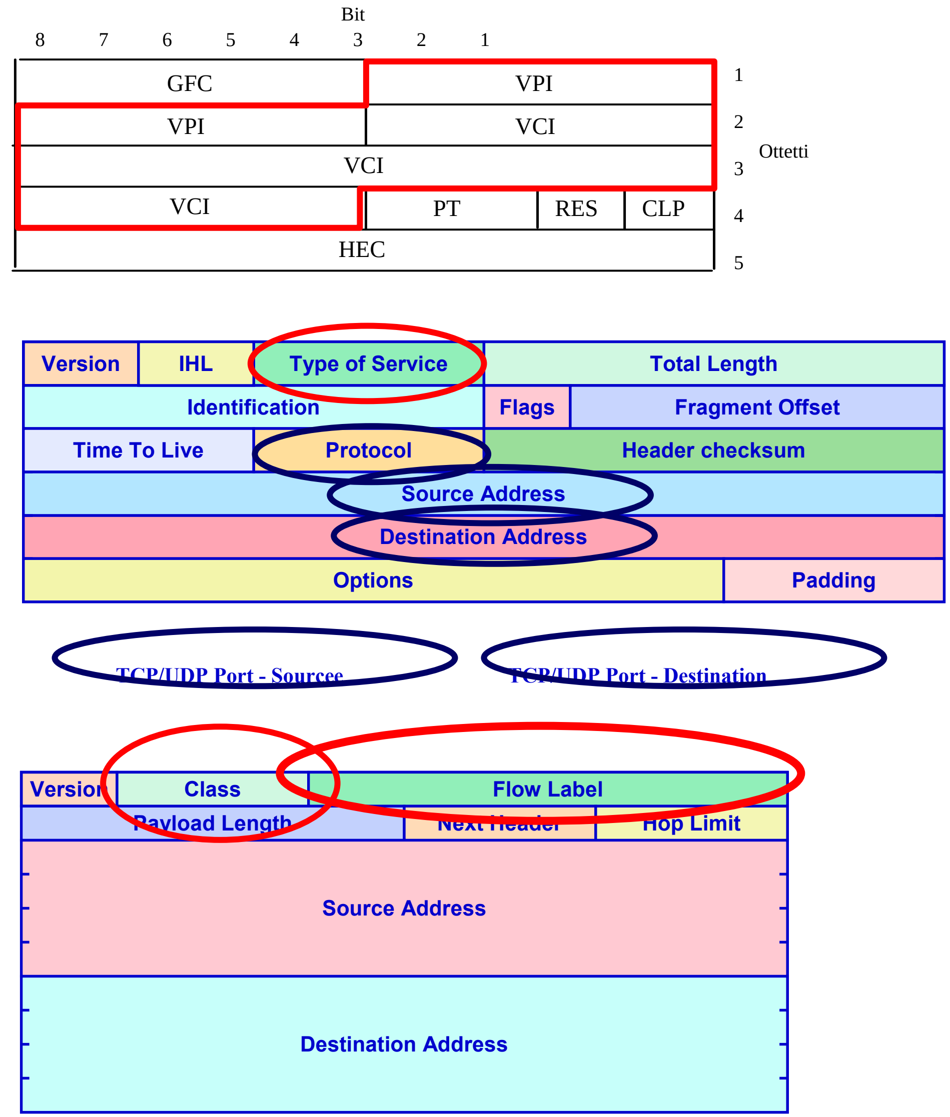

QoS control functions

- Traffic flow identification
- CAC
- Rate control, traffic shaping and filtering
- Bandwidth allocation

- Scheduling (service discipline)
- Flow and congestion control
- QoS Routing

Markov Chains, MDPs, and Optimization

- Discrete Time Markov Chains
- Continuous Time Markov Chains
- Markov Decision Processes
- Dynamic Programming
- Infinite-horizon optimization
- Numerical techniques
- Control Issues in TDM, ATM and IP networks
- Examples

Representation of networks as dynamic systems

- Telecommunication networks are most often modeled as multi-dimensional *complex dynamic stochastic* systems. This means they are interconnected subsystems, whose state depends on time, and whose behavior may be driven by external random variables. In particular, they may be most often represented as *queueing systems*.
- There are many, often equivalent, representations of complex dynamical systems, e.g., in terms of input-output differential equations, transfer function matrices, state equations. One of the most commonly used for queueing systems and networks is in terms of *Markov chains*.

Markov Chains

- Markov chains are Markov processes whose time-dependent random variables (the *state* of the Markov chain) can assume values in a discrete set (the *state space*), either finite or countably infinite.
- The Markov property is essentially a conditional independence of the future evolution on the past (the whole history of the process being summarized in the current state).
- Basically, the chain can be seen as modeling the position of an object in a discrete set of possible locations over time, the next location being chosen at random from a distribution that depends only on the current one.

cnit

Discrete Time Markov Chains

Definition 1.1. A stochastic process $\{X_0, X_1, ..., X_n, ...\}$ at consecutive points of observation 0,1,...,n,... is a DTMC if, for all $n \in N_o, x_n \in S$

$$\Pr\{X_{n+1} = x_{n+1} | X_n = x_n, X_{n-1} = x_{n-1}, ..., X_o = x_o\} = \\ = \Pr\{X_{n+1} = x_{n+1} | X_n = x_n\}$$

▽

Let $S = \{0, 1, 2, ...\}$. The quantities

$$p_{ij} = \Pr\{X_{n+1} = j | X_n = i\} = \Pr\{X_1 = j | X_o = i\}$$

are the one-step transition probabilities of a *homogeneous* chain, i.e., whose conditional pmf is independent of time.

Irreducibility

Definition 1.2. A transition matrix P on the state space *S* is said to be *irreducible* if it is possible for a Markov chain with TPM P to move from any state *i* to any other state *j* in finite time, i.e., if there is a path between any two states in the corresponding transition diagram. A DTMC is *irreducible* if its TPM P is irreducible.

Theorem 1.1. An irreducible DTMC has *at most one invariant distribution* (it certainly has one if it is finite). A DTMC with one invariant distribution is said to be *positive recurrent*.

The invariant distribution measures the fraction of time that the DTMC spends in the various states.

Periodicity and Ergodicity

Let

$$d = \mathrm{GCD}\left\{ n \geq 1 \,\middle|\, (P^n)_{i,i} > 0 \right\}$$

be the greatest common divisor of the number of steps n such that the DTMC can go from state *i* back to itself in n steps (for an irreducible DTMC, d is the same for all states).

Definition 1.3. Let P be an irreducible TPM on *S*. If d > 1, then P is said to be *periodic with period d*. If d = 1, then P is said to be *aperiodic*. ∇

Theorem 1.2. For an *irreducible* and *aperiodic* DTMC with invariant distribution π, the limit $\tilde{\pi}$ exists, is independent of the initial state and coincides with the unique steady-state probability vector. ∇

An irreducible, aperiodic DTMC with all states being positive recurrent is said to be *ergodic*.

Continuous Time Markov Chains

- CTMC's can be viewed as DTMC's with an infinitesimally small time unit. However, a more direct definition can be used.
- To this aim, we recall the properties of an exponentially distributed r.v. τ:

The r.v. τ is exponentially distributed with rate $\lambda > 0$ if

$$\Pr\{\tau > t\} = e^{-\lambda t}, \quad t \geq 0 \; ;$$

If τ is exponentially distributed with rate λ, then

- $E\{\tau\} = 1/\lambda$
- τ is *memoryless*, i.e.,

$$\Pr\{\tau > s + t | \tau > s\} = \Pr\{\tau > t\}, \quad \forall s, t \geq 0$$

Examples

	P	Stationary probability vector(s)	$\tilde{P}$	$\tilde{\pi}$	Unique steady-state probability vector
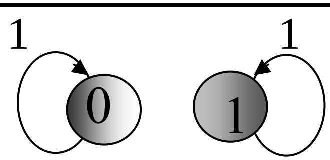	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	Infinitely many	$\tilde{P} = P^n = P, \ \forall n$	$\tilde{\pi} = \pi(0)\tilde{P} = \pi(0)$	None
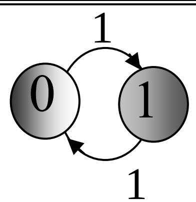	$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	$\pi = [.5 \ \ .5]$	P^n does not converge	$\tilde{\pi}$ does not exist	None
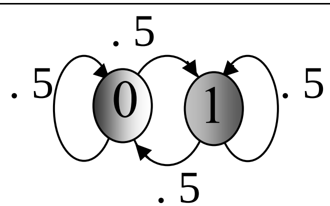	$\begin{bmatrix} .5 & .5 \\ .5 & .5 \end{bmatrix}$	$\pi = [.5 \ \ .5]$	$\tilde{P} = P^n = P, \ \forall n$	$\tilde{\pi} = \pi(0)\tilde{P} = [.5 \ \ .5]$	$\pi = \tilde{\pi} = [.5 \ \ .5]$
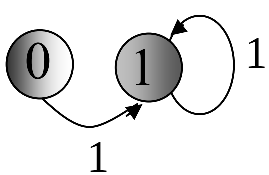	$\begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$	$\pi = [0 \ \ 1]$	$\tilde{P} = P^n = P, \ \forall n$	$\tilde{\pi} = \pi(0)\tilde{P} = [0 \ \ 1] = \pi$	**None that covers the whole state space** $(\pi_0 = 0)$

CTMC (cont'd)

Definition 1.4. Let *S* be a countable set. A rate matrix Q on *S* is a collection $Q = \{q_{ij}, i, j \in S\}$ of real numbers s.t.

$$0 \le q_{ij} < \infty, \ \forall i \ne j \in S, \ \text{and}$$

$$-q_{ii} = q_i \equiv \sum_{j \ne i} q_{ij} < \infty, \ \forall i \in S$$

▽

Definition 1.5. Given a countable set *S*, a rate matrix Q on *S*, and an initial distribution π, the CTMC $X = \{X_t, t \ge 0\}$ is defined as follows

- Choose x_0 with distribution π in *S*;
- If $x_0 = i$, select a random time τ that is exponentially distributed with rate q_i; define X s.t. $X_t = i \text{ for } 0 \le t < \tau$;
- At time t=τ, the process jumps from the initial value i to a new value j, selected independently of τ s.t.

$$\Pr\{X_\tau = j | X_0 = i, \tau\} = \Gamma_{ij} \equiv q_{ij}/q_i \ , \ j \ne i$$

- The construction resumes from there, independently of the process before τ. ▽

CTMC

(cont'd)

- Again, the only information about the trajectory of X up to time t that is useful for predicting the trajectory after time t is the current value X_t.
- Note, in passing, that in processes having non-exponential holding times, but conditionally independent successive jumps (e.g., M/G/1 queueing systems, with Poisson arrivals and general service time distribution), it can be useful to study the (discrete time) Markov chain arising from the process considered at certain jump times (*embedded Markov chain*). Such processes are called *semi-Markov*.
- (It is remarkable that the steady-state distribution of the embedded DTMC in the M/G/1 case is the same as the one of the original non-Markovian process. This is due to the so-called *PASTA* property (*Poisson Arrivals See Time Averages*)).

CTMC (cont'd)

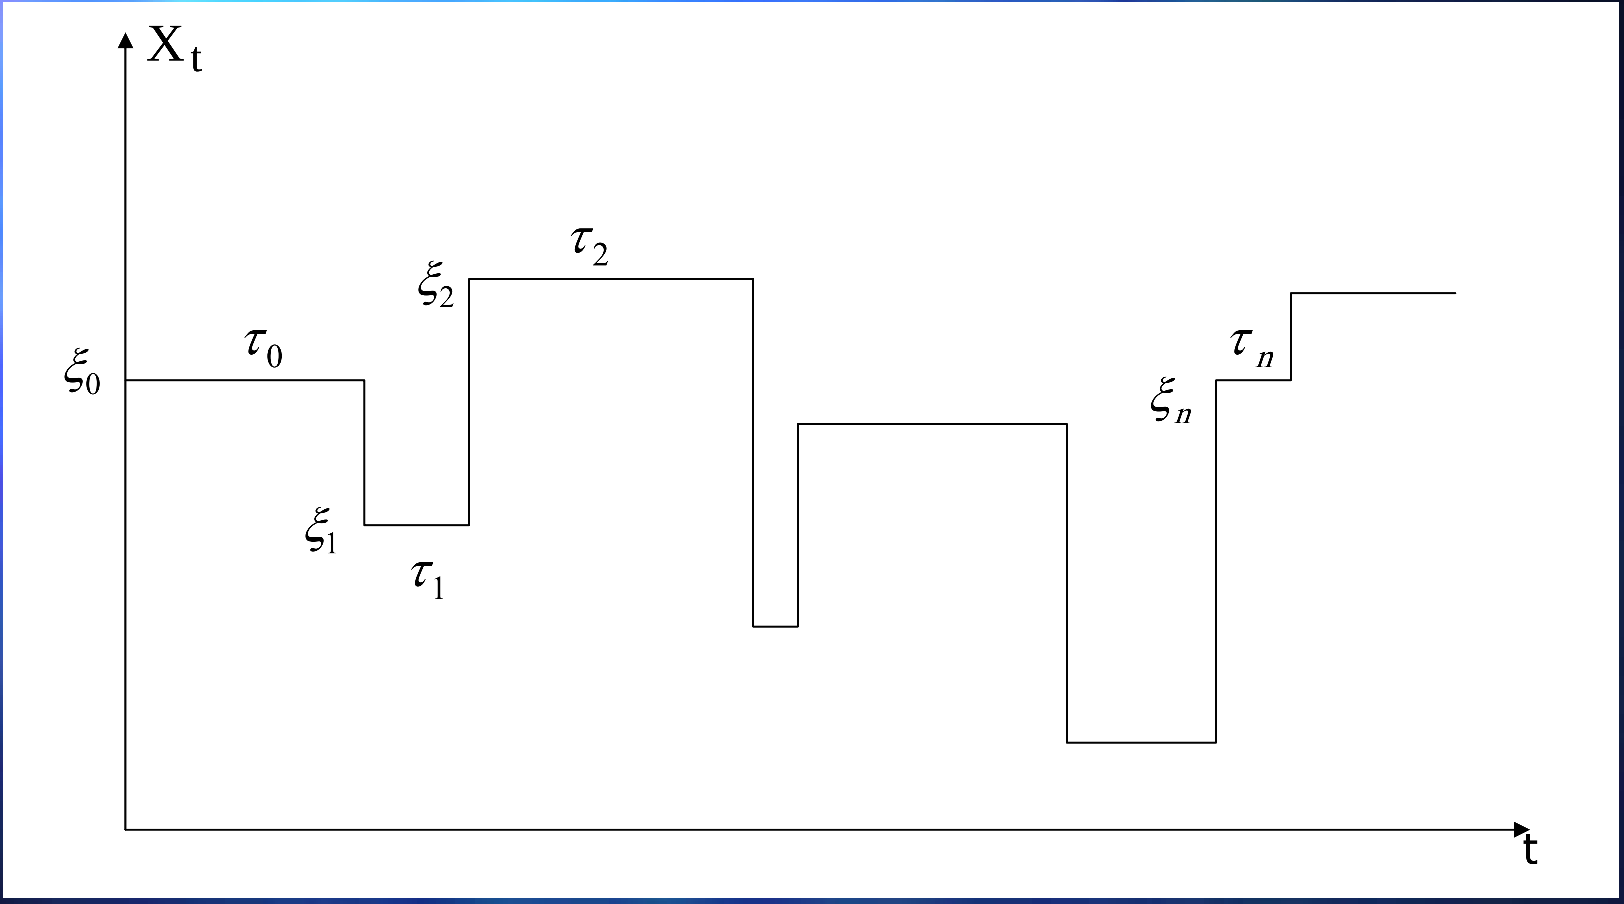

$$\Pr\left\{\xi_0 = i_0, \xi_1 = i_1, \ldots, \xi_n = i_n\right\} = \pi_{i_0} \Gamma_{i_0 i_1} \Gamma_{i_1 i_2} \cdots \Gamma_{i_{n-1} i_n}$$

Controlled Markov Chains

- In both Discrete-Time and Continuous-Time Markov Chains, the elements of the transition or rate matrices may be dependent on a variable u, whose values in a set U(i) determine the transition probability or rate, given state i. We can write $p_{ij}(u)$ or $q_{ij}(u)$ to evidence the functional dependence, or think of the elements of the matrices as being parametrized by u, which represents a control action. We talk in this case of a *controlled Markov chain*.
- In many cases of interest, the sets U(i) may be finite. In general, the action u stems from a *control law* (or *strategy* or *policy*), which determines the action as a function of available information on the process state, either deterministically (*pure policy*) or on the basis of a probability distribution over the action space (*randomized policy*).

Markov Decision Processes

- In general, the information available on the process at time t (either discrete or continuous) may be denoted by I(t), and may represent a whole collection of past observations on the system's state (either perfect or partial); the goal of the control law may be the (functional) minimization of some average cost (or maximization of average revenue), over a time horizon that may be finite or infinite. This is the general setting (in an extended sense) of *Markov Decision Processes (MDP)*.
- A good deal of control problems arising in telecommunication networks (e.g., multiple access, CAC, flow control, dynamic bandwidth allocation among traffic classes) admit a general formulation in terms of MDPs. In some instances, there may even be more than one decisional agent, and such agents may possess different information on the system's state, leading to formulations in terms of game or team theory.

Markov Decision Processes (cont'd)

- As far as finite control horizons are concerned, even problems with partial or imperfect (e.g., noisy) observations of the system's state may be treated efficiently, whenever it is possible to extract from the whole set of past observations (which is growing with time!) a finite-dimensional set of quantities that, loosely speaking, contain all the information in I(t) that is necessary for control purposes. Such a set is called a *sufficient statistic* (or *information state*).
- Over infinite control horizons (that are of interest because they may be characterized by stationary (time-invariant) optimal control laws), little exists regarding the case of imperfect information. *Receding-horizon* approximations (*repetitive control*), either *closed-loop* or *open-loop feedback* may sometimes be viable solutions.
- In order to expose the basic principles, in the following we limit our treatment to the case of perfect state information.

A note on representations of system's dynamics

- It may happen that a networking problem can be formulated more directly in terms of (stochastic, i.e., driven by some *noise* variables) state equations, of the type

$$x_{k+1} = f_k(x_k, u_k, w_k)$$

state

control

exogenous stochastic variable

rather than of Markov chains. As far as discrete time and discrete (finite or infinite countable) state spaces are concerned, it is straightforward to reformulate the dynamics in term of a Transition Probability Matrix.

Control laws

- In the *non-randomized* (*pure*) case, consider control laws of the form

$$u_t = \gamma_t(i), \; i \in S, \; u_t \in U_t(i), \; t = t_0, t_1, \ldots$$

(where t represents a discrete time instant (decision epoch) where the process changes state).

- As regards the *randomized* case, the control law takes on the form

$$\nu_t = \tilde{\gamma}_t(i, u), \; i \in S, u \in U_t(i)$$

s.t. when the process enters state i at time t, action u is chosen with probability $\tilde{\gamma}_t(i, u)$. Obviously,

$$0 \le \tilde{\gamma}_t(i,u) \le 1, \; \forall i \in S, u \in U_t(i) \quad \text{and} \quad \sum_{u \in U_t(i)} \tilde{\gamma}_t(i,u) = 1, \; \forall i \in S$$

Cost (or revenue) functionals

- If a cost is associated to the state and control values, of the type

$$g_t(x_t,u_t) = g_t\left[x_t,\gamma_t(x_t)\right]$$

then the whole cost … en as a sum or an integral over time. We distinguish the cases of discrete- and continuous-time processes, over finite and infinite time horizons, respectively. Let $\tilde{\gamma} = \{\gamma_0, \gamma_1, \dots\}$ represent a policy (a whole set of strategies).

  - *DTMC, finite horizon*:

$$J_{\tilde{\gamma}}(x_0) = E\left\{\sum_{t=0}^{N-1} g_t\left[X_t,\gamma_t(X_t)\right] \Big| x_0\right\}$$

  - *DTMC, infinite horizon*:

$$J_{\tilde{\gamma},\mathrm{disc}}(x_0) = \lim_{N\to\infty} E\left\{\sum_{t=0}^{N-1} \alpha^t g_t\left[X_t,\gamma_t(X_t)\right] \Big| x_0\right\}$$

Discounted cost

$$J_{\tilde{\gamma},\mathrm{av}}(x_0) = \lim_{N\to\infty} \frac{1}{N} E\left\{\sum_{t=0}^{N-1} g_t\left[X_t,\gamma_t(X_t)\right] \Big| x_0\right\}$$

Average expected cost

Cost (or revenue) functionals

- *CTMC, finite horizon*: $$J^{c}_{\tilde{\gamma}}(x_0) = E\left\{ \int_0^{\tau_N} g_t\left[X_t, \gamma_t(X_t)\right] dt \,\middle|\, x_0 \right\}$$

  where τ_N is the instant of the N-th jump of the process. This continuous time cost can be easily discretized over events (by using *uniformization*), to yield a sum over n as in the discrete case (actually, the same is possible, with some more complication, if the final time is a fixed instant T, rather than random).

- *CTMC, infinite horizon*: $$J^{c}_{\tilde{\gamma},av}(x_0) = \lim_{T\to\infty} \frac{1}{T} E\left\{ \int_0^{T} g_t\left[X_t, \gamma_t(X_t)\right] dt \,\middle|\, x_0 \right\}$$

  again, there is an equivalent discrete problem that can be obtained by uniformization. The discounted versions are also possible.

Cost (or revenue) functionals

- In networking problems, the cost (or revenue) is most often associated with:
  - *Loss* of data units (segments, packets, cells, …) in finite buffers
  - *Blocking* of connection or flow requests at the network edge or in the transition across network boundaries
  - *Delay* of data units (in individual buffers or end-to-end)
  - *Throughput* (or *goodput*) at various architectural layers (data link, network, transport, application)
  - *Net user gain* or satisfaction (benefit less price paid for resource utilization)
  - Network or service provider's *revenue*