

Please note

The text in this file has been automatically extracted and may contain minor errors. For the original version please consult the paper copy held in the Swinburne Library.

CHAPTER 2

TOWARDS A CONTROL FRAMEWORK FOR BISDN

2.1 Introduction

This chapter discusses the philosophy behind the control **framework** employed in the later parts of the thesis. In particular, for the control of BISDN we propose: the use of feedback, possibly supplemented by feedforward, control; a hierarchically organised control structure which features both a vertical as well as a horizontal decomposition; that a number of different (dynamic) modelling techniques should be employed; and that the (dynamic) **tradeoff** between service-rate, buffer-space, cell-delay and cell-loss should be exploited.

In section 2.2 we focus on feedback control and define a control horizon, within which feedback controls are effective. We discuss the time and space behaviour of the system, and show that due to the varying, wide ranging "time constant" of the network a structured approach to handle its complexity is required.

In section 2.3 we present such a structured approach: that of a hierarchically organised system. The control functions in the hierarchically organised control system can be allocated to a level related to an appropriate time scale and location. The controls can be calculated in a decentralised fashion with the upper layers coordinating the distributed local units for the overall network benefit. Each control **function** is associated with a certain control horizon, which sets a physical constraint on the location of the control processing relative to the site of the particular physical control event. At lower levels of the hierarchy, control function processing is constrained to be physically close to the control events, so must be distributed throughout the network. At higher levels, the control horizons become so large that physical location of control **function** processing is arbitrary.

In section 2.4 we suggest that different modelling techniques may be **usefully** employed to describe different levels of the decomposed (possibly hierarchically **organised**) control system. Furthermore, as it is increasingly noted, communication networks must have satisfactory dynamic as well as steady-state performance, thus the dynamic aspects of the behaviour cannot be ignored.

The (dynamic) **tradeoff** between service-rate, buffer-space, cell-delay and cell-loss, discussed in section 2.5, allows more flexibility in the effective and efficient control of BISDN.

2.2 System behaviour in time and space

There is a relationship between temporal and spatial distribution of events in a BISDN. A number of connections will generate traffic that is expressed as a distribution of cells over time. The time distribution of the cell stream, as seen at a single point in the network, will also be **influenced** by the spatial distribution of the source nodes of the connections. The influence of spatial distribution on temporal distribution is due to the finite speed of propagation of cells. Of course, delays arise **from** other causes as well as propagation, but all delays may be likened to propagation times and hence to equivalent distances in space (see section 2.5, figures 2.5 and 2.6). In the following discussion, the time scales for the feedback control mechanism and for the network are considered separately. By considering the whole system in this way, we will develop arguments for the distribution in space of certain essential control functions.

2.2.1 Traffic Sources and Services

A BISDN network provides services of various types between users distributed in physical space. The traffic generated by users forms one component of the time and space characteristics of the network. The service types and user demands will be extremely diverse. In general, we need to consider traffic types including Connection Oriented and Connectionless, Continuous Bit-rate and Variable Bit-rate, and with or without a Timing Relation between origin and destination. The Quality of Service (QoS), as perceived by a network user, will depend on the cell-loss, the cell-transfer-delay and the cell-delay-variation (a brief discussion of QoS appears in section 2.5).

2.2.2 Network Physical Resources

The resource **infrastructure**, comprising the physical links and processing nodes, is finite and also distributed in space. The resources that we will consider are of two types: cell waiting places in buffers, and cell-slots on links (the **server resource**). Note that these are substitute resources (see section 2.5, [22], [23])

2.2.3 Feedback Control System

It is well known that feedforward and feedback controls both have essential roles in effective control. Feedforward control is immediate and effective, but only when measurements and behavioural models are accurate. Feedback can be effective even in

the presence of model inaccuracies, but reacts only **after** a disturbance has begun to take effect on the system, and its speed of response is limited by the total delay around the feedback loop. Here we concentrate on the feedback controls, to gain insight into the **influence** of the spatial distribution on the performance of feedback controls.

2.2.3.1 Control Horizon

The *control horizon* will be defined as the shortest time scale possible for a particular control action, and can be usefully visualised in either time or distance units. For example, a control horizon of 100 celltimes at 155 Mbit/sec is equivalent to the propagation delay of a signal on a approximately 60 kms link. The diagram below shows the relationship between the physical link dimensions of a section of metropolitan network and a 100 cell control horizon **circle**^{#1}.

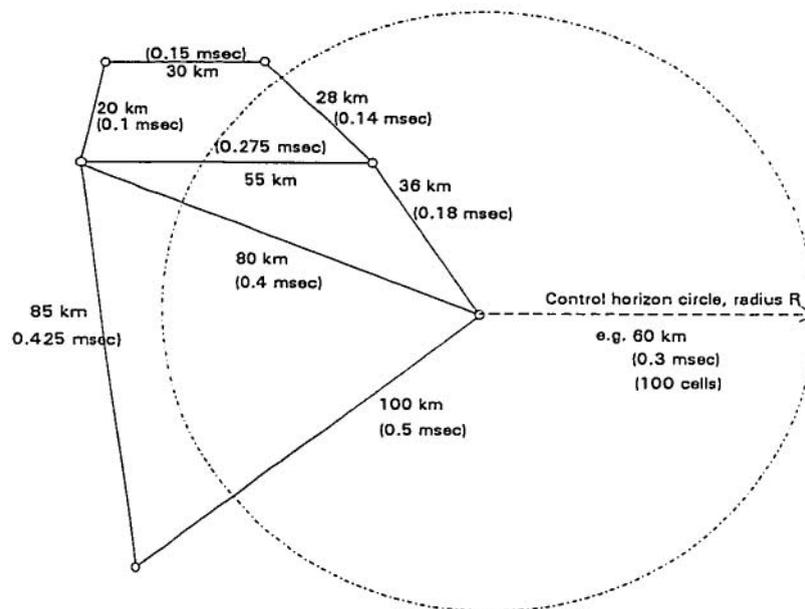


Figure 2.1. Physical and temporal relationship in a network.

The control horizon for a feedback control loop is determined by many factors—these are complex and diverse, but for discussion purposes these will be replaced by the loosely defined idea of "time constant", expressed as the sum of "feedback time constant" and "network time constant" (see figure 2.2). Note the use of-"time constant" here is different **from** the classical, rigorously defined, time constant.

^{#1} Note that for averaging periods of say 600 cells (a reasonable averaging period for flow measurements), corresponding to 360 kms (a control horizon that covers most metropolitan networks), propagation delay may be an insignificant constraint on feedback control.

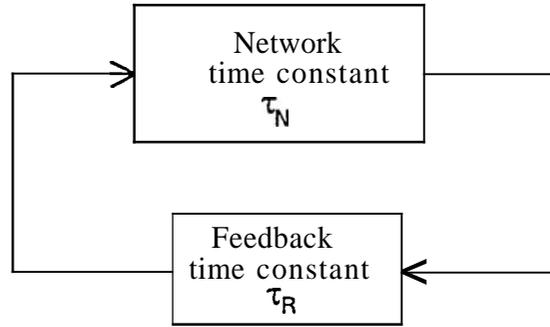


Figure 2.2. Factors influencing the control horizon.

The "feedback time constant" is a **function** of the measurement interval, the propagation delay **from** point of measurement to point of control action, and the processing time required to calculate the new control actions.

$$\tau_R = f(\tau_m, \tau_p, \tau_{pr}) \quad (2.1)$$

where:

τ_m is the measurement interval, i.e. the time involved in making a measurement,

τ_p is the propagation delay, and

τ_{pr} is the processing time (which includes the time required to make the measurements and calculate the feedback signals, e.g. unfinished work distribution in a buffer).

The "network time constant" depends on the duration of the transients following a control action. This is more difficult to represent, since it depends on a number of network and user attributes and requirements. Among others it depends on:

$$\tau_N = f(\lambda_{call}, \lambda_{call}^d, \tau_h, \frac{\partial \lambda_{call}}{\partial t}, D_{max}, P_{loss}^{max}, \lambda_{cell}, \frac{\partial \lambda_{cell}}{\partial t}, C_{link}, B_{link}) \quad (2.2)$$

where:

λ_{call} is an $S \times 1$ vector of the call arrival rate

λ_{call}^d is an $S \times 1$ vector of the call death rate

τ_h is an $S \times 1$ vector of the call holding time

$\frac{\partial \lambda_{call}}{\partial t}$ is an $S \times 1$ vector of the rate of change of the arrival or death rate of the call rate

D_{max} is an $S \times 1$ vector of the maximum tolerable delay

P_{loss}^{max} is an $S \times 1$ vector of the maximum tolerable cell-loss

λ_{cell} is an $S \times 1$ vector of the cell flow rate

$\frac{\partial \lambda_{cell}}{\partial t}$ is an $S \times 1$ vector of the rate of change of the cell arrival rate

C_{link} is the link service-rate,

B_{link} is the link buffer-space (if logically divided among the S type calls, then its individual components must be taken into account).

S number of calls (or the number of different types of calls; assuming that the properties of a call can be adequately represented by a general class i , where $i = 1, \dots, S$).

It is commonly assumed that network behaviour can be separated into different components, each of which has its own "network time constant". These components can then be grouped into broad categories by time constant. At least five broad time constants (levels) have been proposed in the literature [24], [25], mainly using intuition: cell level (tens of microseconds); burst level (tens of milliseconds to seconds); call level (tens of milliseconds to tens of minutes); virtual path level (several seconds to tens of minutes); and the network level (tens of minutes, hours, days). The recognition of the existence of different time scales for different controlled events has been an **important** first step in providing effective controls. What is required, in addition to the above decomposition of the time constant, is an association of the control functions for each broad level with their spatial distribution—that is the above time decomposition provides only the vertical level decomposition. We will now consider the other important point of view—the horizontal decomposition—by examining some of the functions that need to be performed in the network.

2.2.3.2 Control Functions

Some examples of network control functions are presented in table 2.1. For each control **function** the table shows: the probable time scale and location; the controlled event and its location; and the main factors affecting the time response. Note that the probable time scale and the probable physical location of the control function are deduced from the type of information the control functions require and also the frequency of updating of the control output (as proposed in the literature, e.g. see [26]).

| Control function (examples) | Controlled event | Controlled event location | Probable time scale level of the control function | Probable physical location of the control function | Main factors affecting the time response |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|---------------------------------------------------|-----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> - reconfiguration of the network (e.g. creation and deletion of VPs) - resource reservation (e.g. cell-rate for VPs) | <ul style="list-style-type: none"> - network topology, VP topology - service-rate | spans all nodes of the network | network level | main control centre | $\lambda'_{call}, \lambda^{d,i}_{call}, \tau'_h,$ $\frac{\partial \lambda'_{call}}{\partial t}$ at |
| <ul style="list-style-type: none"> - routing (e.g. selection of alternative VPs) - dynamic service-rate control - flow control <ul style="list-style-type: none"> -rate based -window based | <ul style="list-style-type: none"> - VP connection, Virtual call connection - service-rate - cell-flow | spans several nodes from origin to destination | VP level virtual connection level | main control centre or Origin node or nodes along the connection path | $\lambda'_{call}, \lambda^{d,i}_{call}, \tau'_h,$ $\frac{\partial \lambda'_{call}}{\partial t}$ $D'_{max}, P^{i,max}_{loss}, \lambda'_{cell},$ $\frac{\partial \lambda'_{cell}}{\partial t}, C_{link}, B_{link}$ |
| <ul style="list-style-type: none"> - Call Admission Control (CAC) - dynamic service-rate control - flow control - fast reservation protocol -explicit congestion notification | <ul style="list-style-type: none"> - call, connection - service-rate - cell-flow - service-rate - cell-flow | spans several nodes from origin to destination | call level | main control or Origin node or nodes along the connection path | $\lambda'_{call}, \lambda^{d,i}_{call}, \tau'_h,$ $\frac{\partial \lambda'_{call}}{\partial t}$ $D'_{max}, P^{i,max}_{loss}, \lambda'_{cell},$ $\frac{\partial \lambda'_{cell}}{\partial t}, C_{link}, B_{link}$ |
| <ul style="list-style-type: none"> - Cell discarding - priority control - usage parameter control - traffic shaping - link server discipline - link buffer | <ul style="list-style-type: none"> - cell - cell - cell-flow - cell-flow - service-rate - cell-space | node | cell level | nodes along the connection path | $D'_{max}, P^{i,max}_{loss}, \lambda'_{cell}$ $\frac{\partial \lambda'_{cell}}{\partial t}, C_{link}, B_{link}$ |

Table 2.1. Examples of control functions showing: the controlled events and their location; probable time scale levels and location of functions; and factors affecting the time response.

From table 2.1 we get a glimpse of the variation of the time constant as well as the geographic separation of the events. The time constant can vary over several orders of magnitude. Even within the vertical levels identified in the literature—for example the call holding time for a video connection may be tens of minutes whereas for a file transfer it can be tens of milliseconds. The geographic separation can span all the nodes along a **connection—from** a few kilometres to thousands of kilometres.

For simplicity of exposition we focus on one controlled event, that of the service-rate. From table 2.1 we observe that it is associated with all four levels, indicating that elements of its time constant extend over all the identified levels. Especially, we note that the time constant can vary widely with the instantaneous connection mix, and can be time varying, even within an identified level. However in the literature the vast majority of the schemes discussing **service-rate** control focus on one level in isolation (possibly using the assumption that the time scales associated with the various levels are widely separated and hence no interference is experienced between the levels) and propose solutions that do not take into consideration the interactions between the different levels and the wide variations of the time constant values (even within the identified levels). The dynamic behaviour of the event is largely ignored.

Additionally, as can be seen **from** table 2.1, the controlled events (e.g. service-rate) are geographically distributed, and due to the limit imposed by the control horizon, cannot be effectively controlled unless the control processing is also distributed in space.

Based on the above discussion, we are prompted to find an approach that will allow us to tackle the diverse range of the time constants as well as the geographic distribution of controlled events. Only by decomposing events both vertically as well as horizontally (in time and space), and organising these control tasks in an orderly fashion will we be able to design coherent and effective controls.

2.3 Hierarchical (multilevel/multilayer) control

A hierarchical multi-level/multi-layer control structure (Mesarovic et al [27], Singh and Titli [28], Siljak and Sundareshan [29], Cruz [30], Findeisen et al [31], Mahmoud et al [32], Haimes et al [33]) is investigated here for the effective and efficient control of this large scale complex network system.

There is an inherent constraint on a control system, irrespective of the control structure it is placed in. All elements of a (feedback) control loop for a given function must be within the physical constraint of the control horizon circle for that function. This sets a constraint on situations in which feedback control is feasible. Hierarchical control normally is used to tackle the problem of complexity. Decomposition allows the

complexity of each block to be contained. An inherent consequence is usually to progressively relax the time scales at progressively higher (vertical) levels of the hierarchy. Hierarchical control is also relevant to the problem of physical distribution of the network and the finite propagation time between elements. Decomposition (horizontally) may allow lower level blocks to be limited to sets of physically close elements. The inherent relaxation of time scale at high levels of a hierarchical system can be directly associated with a relaxation of physical proximity constraints on the control elements of that **functional** level of the hierarchy. Therefore, within a hierarchically organised system, control **functions** can be allocated to a level related to an appropriate time scale and location, and controls can be calculated in a decentralised fashion with the upper layers coordinating the distributed local units for the overall network benefit.

Our proposed structure is intended to provide adaptability and robustness. **Siljak** [34] has shown that a complex system, when synthesised of interconnected stable subsystems, has highly reliable stability properties. A locally stabilised system is robust and has a high tolerance to **nonlinearities** in the interactions among the subsystems. The modularisation into simple well defined tasks also allows flexibility as well as a natural evolution path as technology changes.

The **key features** of hierarchical control are:

- Conceptual simplification is achieved by the decomposition of the system into several subproblems organised at different levels.
- Decomposition addresses complexity by yielding small subsystems of low dimensionality, and as a result, a reduction of the dimensionality of the overall system.
- The subsystems can be spatially distributed, and have limited communication with one another.
- Subsystems make decisions autonomously using private information. This allows fast acting local controllers to be implemented, coordinated by higher levels toward achieving global goals. Organisation of the control problem in decision hierarchies can be based on aggregated models, with different hierarchies solving different decision problems (the lower levels solve the more detailed problems whereas the highest level solves the global problem based on a more aggregated model of the overall system).
- Each subsystem can be described by a different class model. These models can be linear or nonlinear, static or dynamic or less traditional, such as Discrete Event

System (**DES**) models [35]. Therefore appropriate models can be used in order to adequately describe the behaviour of each component of the decomposed system.

- Each subsystem can be optimised independently using optimisation techniques appropriate for the particular subsystem. Note that this independence is provided by the fixing of certain variables in the optimisation and control problems for the subsystem. These values are updated by a higher level using coordination variables (for a discussion on decomposition and coordination see appendix 4.1, page 159).

Note though that there is a cost, delay, or distortion in transmitting information between the levels in the hierarchy.

Historical note

A number of papers have appeared in the open literature discussing hierarchical structures for the control of communication networks, some examples are: a general hierarchical structure for the control of BISDN (Pitsillides et al [36], [37]); a multidimensional framework centered around a rate-based access control scheme (Ramamurthy et al [38]); a general multistrata framework (the **M-architecture**) for resource allocation in which a stratum is defined both by the layer of traffic flow and the controlled level, and the resources can be allocated to cells, bursts, calls and flows (Filipiak [25]); a stratified network management scheme (Warfield et al [39]); a general layered approach for network management and control with the details of particular functions (task and location) left to the users (Campbell and Everitt [40]); a layered and distributed congestion-control framework (Eckberg et al [41]); a two level hierarchical structure for routing and flow control (Muralidhar and Sundareshan [42]); a two level hierarchical structure for call admission and bandwidth allocation at a single ATM channel (Bola et al [43]); a two layer congestion control scheme (Ren and Meditch [44]); a three level resource management architecture based on a preventative congestion control strategy (Anido, Bradlow, and co-workers [45]); bandwidth allocation in three levels, the packet, burst and cell level (Hui [24]); and Jain [14] who argued on strong intuitive grounds the assertion that no single congestion control strategy is adequate implying the need for more than one level. With a few exceptions, e.g. Warfield et al, Eckberg et al, Pitsillides et al, the proposed schemes have been motivated by the implementation of one solution to the problem, organised in a multilevel structure.

Our work differs **from** published work of others in the sense that we motivate a general hierarchical organisational and control structure **from** an analysis of the network behaviour both in **time and space**.

2.3.1 The proposed hierarchical control structure

We consider three fundamental classes of control unit (see figure 2.3):

Local Unit, LU:

At the lowest level a Local Unit (LU) uses local measurements, takes direct action, and may be coordinated by a higher level unit. Because of the physical separation of the local units and their limited control horizon, local units must have the ability to act autonomously.

Supremal Unit, SU:

At intermediate levels Supremal Units (SU) coordinate groups of lower level units. The intermediate SU level may need to communicate with higher levels as well as lower levels so that achievement of local objectives contributes to the global network benefit.

Overall Supremal Unit, OS:

At the highest level the Overall Supremal unit (OS) provides global coordination for overall network benefit. A number of conflicting and possibly noncommensurable objectives can be formulated at the overall supremal (OS) level, for example minimisation of total network delay, **minimisation** of total network loss, maximisation of bandwidth utilisation and **maximisation** of total network revenue. The overall supremal unit is responsible for resolving this conflict and directing the behaviour of lower level units.

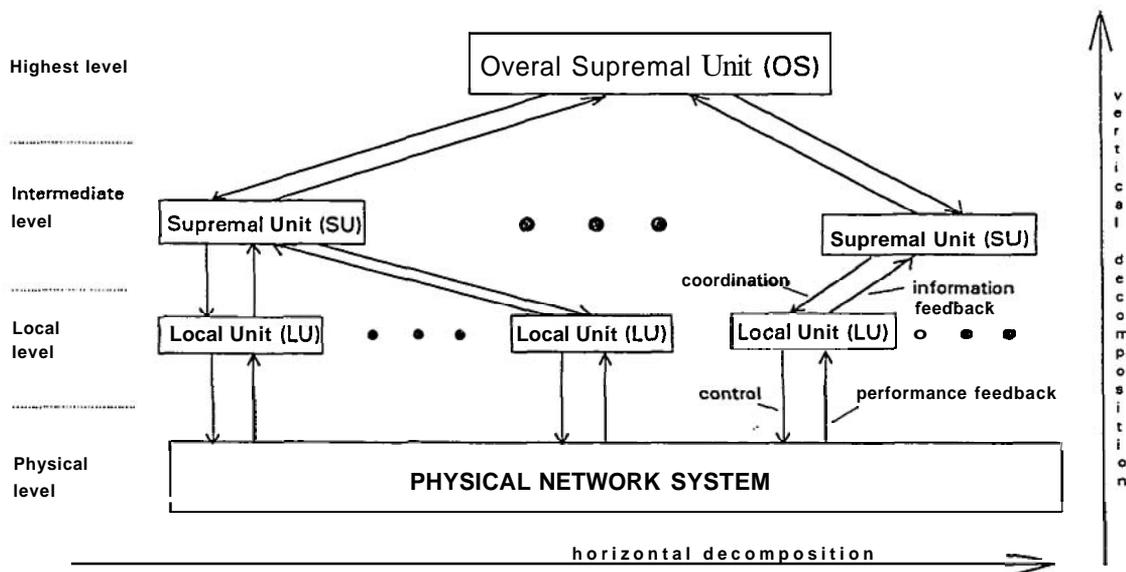


Figure 2.3. A hierarchically (multi-level/multi-layer) organised control structure.

Note that there may be more than one level (with possibly their own local units) within each broad level of the hierarchy.

As an example we show a particular implementation of a BISDN control structure, using these general concepts, in figure 2.4. This example is mainly concerned with service-rate control and it makes use of the Virtual Path (VP) concept (defined as a pre-established route through the network into which Virtual Channels (VCs) can be grouped [46], [47], [48], [49], [50]). Note that the VP is a convenient basis for dynamic service-rate control since it represents flows of traffic (rather than individual connections). At the local level, the link LUs use local feedback, under some higher level direction, to allocate the time of the server resource (the link) amongst the VPs which share the resource. At higher levels of the hierarchy SUs are responsible for VP traffic control and resource management tasks. The function of the SUs is to allocate appropriate resources among the competing VPs, select alternative paths, shape traffic at the source, and accept/reject new call attempts. These functions can be organised into several levels, three of which are shown in this example: SUL3, responsible for coordinating the actions of all Origin-Destination (OD) pairs originating at a node; SUL2, responsible for one OD pair with (possibly) multiple VP paths through the network; and SUL1, responsible for one VP path of the OD pair. The likely location of the SUs is at the origin node of each OD pair. At the highest level of the hierarchy the OS is responsible for directing the behaviour of lower level units toward achieving the global objective(s), for example the maximisation of the total bandwidth utilisation. The likely location of the OS is at the main control centre.

2.3.2 Time scales and decompositions

As already discussed, there is an inherent relationship between vertical and horizontal decompositions. It is this relationship that will determine the necessary extent of the decompositions. Intuitively, one can state that: if the control horizon is not long enough to cover the span of the controlled event then further decompositions are necessary. Note that the investigation of this relationship is beyond the scope of this thesis: its formal study is recommended.

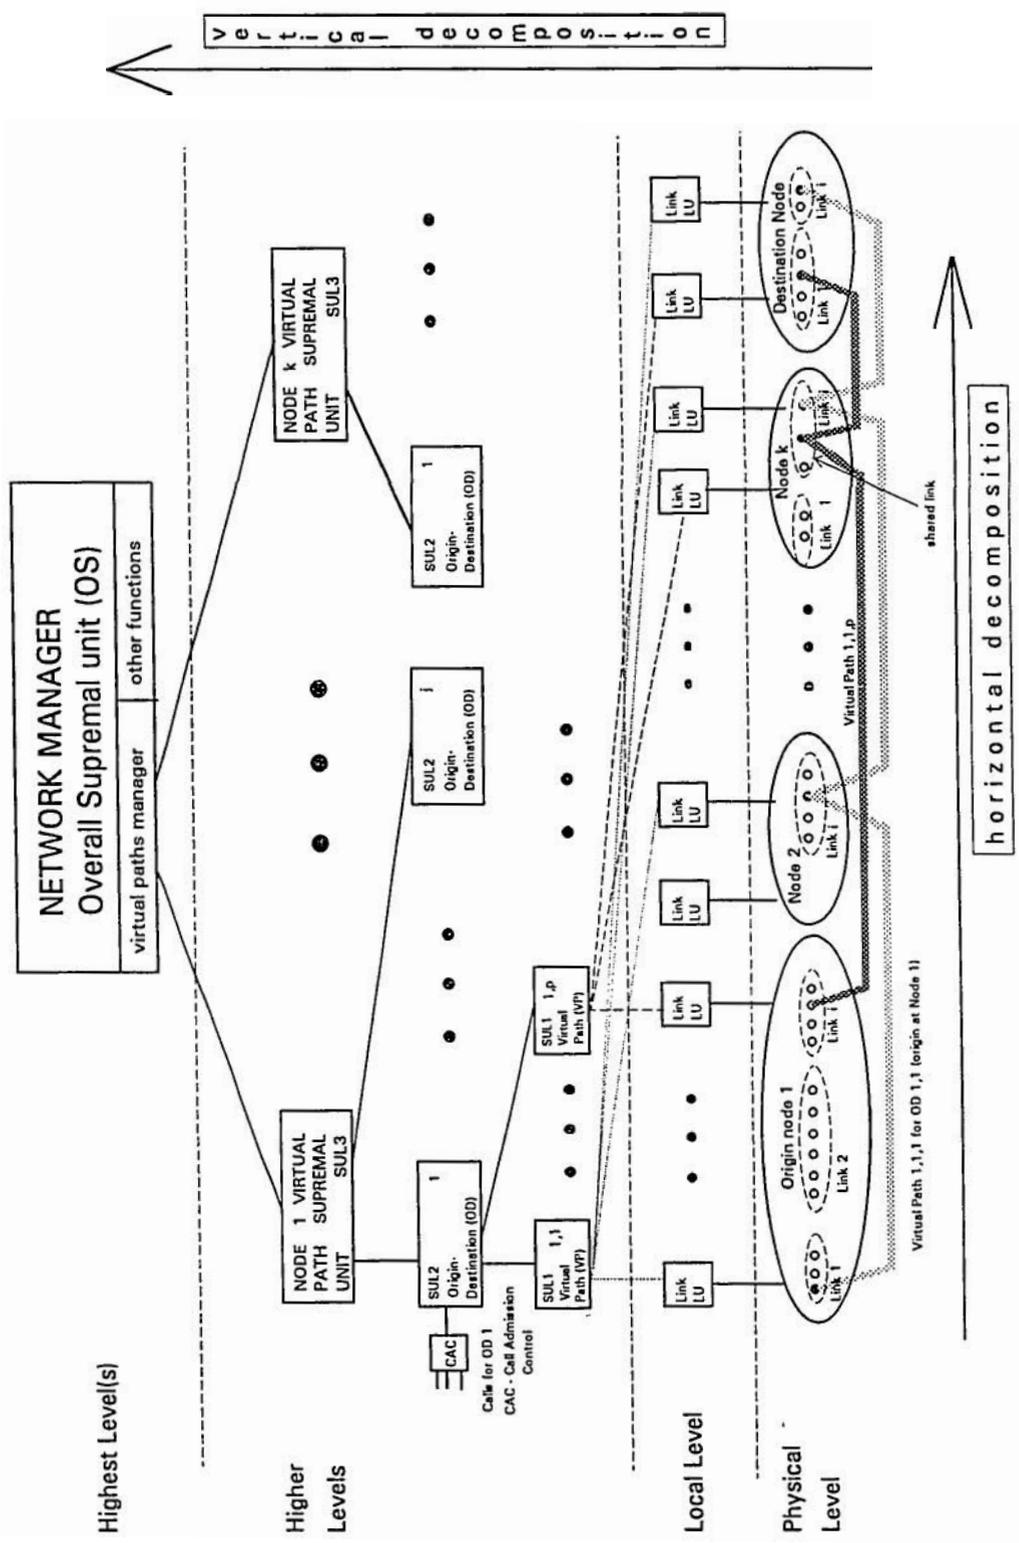


Figure 2.4. Example of a hierarchically (multi-level/multi-layer) organised control structure.

2.4 Modelling techniques for the dynamic control of BISDN

In this section, we look at modelling techniques suitable for the dynamic control of BISDN. The problem of modelling is inherently difficult due to the complexity of the system. In order to make analysis more tractable, the system must be modelled as an interconnection of simpler sub-systems, each giving adequate "lumped average" description of the total behaviour of the system at that level. The **fundamental** problem of modelling is that accurate characterisation of this system requires an excessively large state space. For example, assuming it can be modelled by a Discrete-State Continuous-Time (DSCT) Markov process, a 3 node network system with finite (100 cell places) buffers requires at least 10^{12} discrete states (see section 4.4, page 129). This precludes any hope of deriving a simple on-line controller by this method which is computationally feasible only for the simplest of cases. The key **simplifying** insight is that instead of dealing with "true" states of the complex system, we can work with aggregated quantities such as cells in a queue, flow of cells into a queue, and so on. The advantage of this approach is that

- expected values are real numbers as opposed to states which are integers, so we are dealing with continuously variable quantities instead of discrete-state variables,
- the model has lower state space dimensions, and therefore leads to algorithms which have reasonable computation time.

Something is lost of course: the innumerable transitions among the discrete states are lumped together as additive "noise".

key assertions:

- the most detailed model of DSCT Markov process model is not tractable, hence
- the use of expected values of quantities, which can be measured. Also
- the dynamic aspects of BISDN cannot be ignored, and that
- the system may be modelled as an interconnection of simpler sub-systems (vertical and horizontal decomposition), giving adequate "lumped average" description of the total system behaviour at each level.

Different modelling approaches are therefore required at different levels. In this thesis we make use of a range of models (e.g. in Chapter 3 we use a stochastic difference

equation, in Chapter 4 a dynamic fluid flow model and in Chapter 5 a probabilistic model).

2.5 The cell-delay, cell-loss, buffer-space and service-rate performance and resource tradeoffs

The **Quality of Service (QoS)**, as perceived by a network user, will depend on cell-loss, cell-transfer-delay and cell-delay-variation. The **QoS** target depends on the requirements of the individual connection (or connection mix). The achievement of the target **QoS** is determined by the service-rate and buffer-space resources allocated to the connection (or connection mix).

2.5.1 Cell-transfer-delay and Cell-delay-variation (CDV)

Queuing and propagation delays may be put into a users perspective as follows: A typical cell-transfer-delay encountered by a cell passing from an origin to destination can be seen in figure 2.5 for a 100 kilometre link, 1000 cell queue, and a service rate of 155 Mbit/sec. For voice, the packetisation delay dominates, and total delay is equivalent to less than 5 metres propagation at the speed of sound. For video telephony, 5 Mbit/sec video (assuming no frame delay) and 10 Mbit/sec data, transfer delay is dominated by queuing delay, but this is still negligibly small from a user perspective. [Note: packetisation delay, the delay experienced in "filling" up a cell, is taken as the quotient of the length of the information field of a cell (in bits) and the connection bit rate; and depacketisation delay (dejittering and reassembly) is made equal to the maximum queuing delay—this allows for the worst cell delay variation in the arrival of cells].

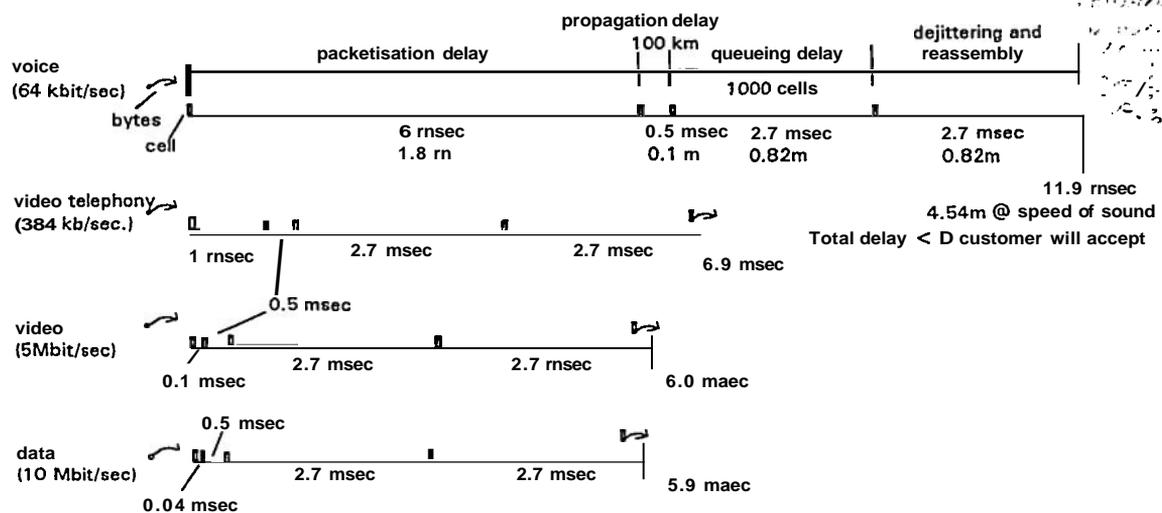


Figure 2.5. Time delays for various traffic types over 100 kms link and 1000 cell queue.

In figure 2.6 the voice example is extended to a 1000 kms link, with queueing delays of 1000 and 10000 cells. Voice is considered since it has a low cell-rate and therefore it features high packetisation and depacketisation delay.

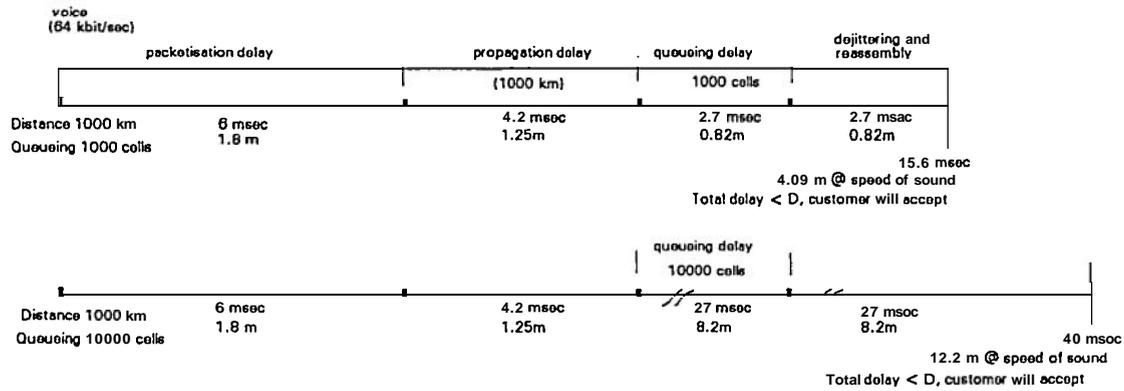


Figure 2.6. Time delays for voice over 1000 kms link and 1000 and 10000 cell queues.

It can be seen that from a customer point of view, fairly large distances (for a 1000 kms link the propagation delay is 4.2 msec) and/or large queues (the queueing delay for a 1000 cell queue is 2.7 msec) can be tolerated. For example, according to CCITT recommendation G.114 [51] the maximum permissible one-way delay for telephone connections is 400 msec^{#1} (although echo cancellation must be employed when a one-way delay of about 25 msec is exceeded). The de-jittering and reassembly process imposes only a very small demand in terms of the storage memory requirements. For example, de-jittering and reassembly for a voice connection requires a cell buffer of

$$b_d = \left\lceil \frac{\tau_q \times c^{dest}}{C^{link}} \right\rceil \tag{2.3}$$

where

b_d is the size of the de-jittering and reassembly cell buffer,

τ_q is the queueing delay in number of cells,

c^{dest} is the service-rate of the destination terminal (in most cases the same as the transmitting terminal),

$\lceil x \rceil$ is the next highest integer toward $+\infty$ (round up).

Example: for a 10000 cell-queueing-delay experienced by a voice connection, transmitted at a rate of 64 Kbit/sec, we require a 5 cell-buffer (i.e. 240 bytes) for de-jittering and reassembly.

^{#1} This delay figure takes into account connections involving satellite links.

Note that certain applications, for example some data file transfers, are delay tolerant; considerably more than what has been shown for voice.

2.5.2 Cell-loss

Acceptable cell-loss performance is dependant on the tolerance of a connection to **information** loss. Quantitative assessments of the tolerance to loss for different service types are hard to find in the literature. A typical figure quoted in the literature as an acceptable target for BISDN is a probability of loss of $P_{loss} = 10^{-9}$ but this seems a much lower rate than necessary for many services. Voice is reported to be more tolerant to losses than the above figure, and other traffic types may also have a higher tolerance to cell-loss. Different targets for different traffic types are more relevant to customer requirements, but the network controls may then become more complex. Its worth mentioning that CCITT suggests two classes of cell loss rate (on a queue by queue basis): 10^{-6} , for example for voice connections; and 10^{-8} or 10^{-10} , for example for data. In Chapter 3, it is shown that the probability of loss can be regulated, by using a control theoretic approach.

2.5.3 The tradeoffs

Cell buffer capacity and service-rate have been shown to be substitute resources in a queueing network [22], [23]. We consider their effect on cell-delay and probability of cell-loss.

If we allow the buffer size to tend to infinity, then the probability of loss, for a finite input rate will tend to zero, i.e. if

$$b \rightarrow \infty \Rightarrow P_{loss} \rightarrow 0, \quad \lambda < \infty \quad (2.4)$$

where

b is the number of cell-waiting places in the cell buffer,

λ is the input rate into the queueing system, and

P_{loss} is the probability of cell-loss (from the cell buffer).

Note that for system stability, we also require the input rate into the system to be less than or equal to the link server service-rate i.e. $\lambda \leq C^{link}$.

Since maximum delay is directly proportional to the maximum number of cells in the queueing system (state of the cell buffer \mathbf{x}), for a maximum tolerable delay and a finite cell-buffer size the above postulate must be modified to

$$x \rightarrow \text{Min}(x^{\max}, m) \Rightarrow P_{loss} \rightarrow \text{Max}(P_{loss}^{\max}, 0), \quad x^{\max} \in X, \quad \lambda \leq C^{\text{link}} \quad (2.5)$$

where

x is the cell-buffer-state (and hence delay),

x^{\max} is the maximum cell-buffer-state (and hence maximum delay),

D_{\max} is the maximum tolerable delay,

P_{loss}^{\max} is the maximum tolerable probability of cell-loss (from the cell-buffer),

X is the set of maximum cell-buffer-places (mapped from the set of maximum tolerable delays $\{D_{\max}\}$), and

C^{link} is the link server service-rate.

A demonstration of this tradeoff can be seen in table III of Heyman et al [21]. A real video sequence of half an hour duration is used as the input to a multiplexer with a finite cell-buffer and a fixed service-rate. The experiment is repeated for 5 different cell-buffer sizes determined by the constraint on the maximum allowed delay (for example, for a link rate of 45 Mbit/second, a cell size of 64 octets and a maximum allowed delay of 5 msec the buffer size is equal to 439 cells). The observed Probability of cell-loss ($\times 10^{-6}$) versus delay is shown in the figure 2.7 below:

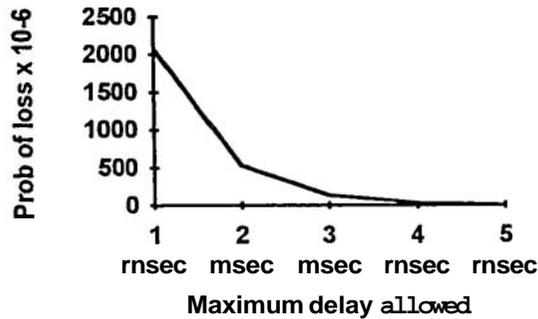


Figure 2.7. Tradeoff between cell-delay and cell-loss.

Observe the improvement in the probability of cell-loss by nearly 3 orders of magnitude (from 2.070×10^{-3} to 2.88×10^{-6}) for what appears to be a modest increase in the maximum tolerable cell-delay (from 1 to 5 msec).

The discussion so far assumes that the link service-rate is constant, but if it is allowed to also vary then we have a means of influencing both the cell-loss and cell-delay, i.e.

$$C^{\text{link}} \rightarrow \infty \Rightarrow \{P_{loss} \rightarrow 0, x \rightarrow 0\}, \quad \lambda \leq C^{\text{link}}. \quad (2.6)$$

Thus, from the above discussion, we can observe the **tradeoff** between service-rate, buffer-space, cell-delay and cell-loss. Also observe that cell waiting places in a buffer and service-rate are substitute resources. As long as the **QoS** constraints on cell-delay and cell-loss are not violated, there is scope for the service-rate and buffer-space **tradeoff**. Note that more substantial gains can be made if the **tradeoff** is also dynamic.

Therefore the (dynamic) **tradeoff** between cell-delay, cell-loss, buffer-space and **service-rate** must be exploited in order to gain network efficiencies for both the customer and the network operator. For example, in Chapter 3 we offer a control scheme with a regulated **QoS** measure, within which these tradeoffs can be exploited by appropriate setting of the reference value.

Additionally, note that the **tradeoff** between buffer-space and control horizon range cannot be ignored (since by providing for more buffering, within the delay tolerance constraint, the control horizon can be extended).

2.6 Conclusions

The main conclusions are:

- the use of feedback for the effective control of BISDN is feasible despite the propagation delays, as long as it is constrained to lie within the control horizon (as for example by appropriate vertical and horizontal decompositions);
- a hierarchically organised control structure (featuring both a vertical as well as a horizontal decomposition) is necessary for the effective control of BISDN;
- different (dynamic) modelling techniques must be employed for the control of BISDN;
- the (dynamic) **tradeoff** between service-rate, buffer-space, cell-delay and cell-loss cannot be ignored.

The results in this chapter motivate our general approach to control in BISDN. For example, in Chapter 3 we integrate the problem formulation of CAC and flow control (we use a stochastic difference equation to describe the system, thus taking into account its dynamic behaviour), with the essential coordination handles provided in the problem formulation. In Chapter 4 we solve the dynamic service-rate control problem at the VP level (making use of a dynamic fluid flow type equation), and finally in Chapter 5 we provide an example of a hierarchically organised solution to the service-rate control problem (at the higher level a constrained optimisation problem using a steady state performance measure—derived from probabilistic principles—is used).