

# Emergent properties of the BT SDH network

J Spencer, D Johnson, A Hastie and L Sacks

---

*When planning an SDH network, explicit structure, such as rings or a network hierarchy, is often imposed to allow for easier network protection and management. Decisions on node connectivity are also heavily dependent on available transmission capacity and network geography, as well as the demands placed on the SDH transport layer. The strictly imposed structure therefore makes it unlikely for unplanned properties to appear but here we describe how the BT SDH network exhibits emergent power-law properties in a range of metrics. These properties are similar to those previously found in the Internet, but the Internet in contrast is not globally planned, has adaptive elements such as dynamic routing and peering agreements, lacks explicit imposed structure, and is less coupled to transmission topologies.*

*This paper shows that even with a wide range of restrictions and controls the BT SDH network topology follows power-laws and we offer possible sources for them, concentrating on the possible effect of adjacent network layers. The existence of these properties has wide implications in network modelling, as well as network scaling, growth, and robustness analysis.*

---

## 1. Introduction

When modelling synchronous digital hierarchy (SDH) networks it is usual to consider the algorithms responsible for the planning of these networks. These algorithms are well defined and consider a wide range of inputs describing the network scenario, including the demands of the network, the configuration of nodes, available physical layer capacity as well as restrictions placed on the design by various technological issues. During planning it is also common to incorporate deliberate structures such as rings to achieve additional attributes such as resilience. Hierarchies are also used in the network to make the management of capacity easier.

It would seem that the real-world result could therefore be easily deduced from knowledge of the demands and the planning algorithm used. In this paper we will show that even with strict planning, emergent and unplanned topological traits appear in the real world SDH network that we consider. The existence of such traits within large-scale networks is not without precedence, however. It has been previously shown [1] that the Internet demonstrates similar traits at a number of levels, from router connectivity to the connectivity of sub-domains. In stark contrast, however, the Internet is not globally planned like an SDH network, does not have enforced structure and often uses dynamic routing techniques to adapt to network conditions. It may not therefore be surprising to find such emergent topological traits there.

In this paper we begin by examining the overall telecommunications system and its constituent layers. We

then briefly describe the macroscopic traits found in the Internet, an example of one of the layers, the Internet protocol (IP) layer, and then demonstrate the existence of these traits in a deployed SDH network. We describe additional traits to those found in the Internet, including some describing topology, bandwidth distribution and geographic connectivity. We then attempt to offer possible explanations to the source of these traits and describe some existing models that attempt to emulate these traits in the Internet and their applicability to SDH networks.

## 2. Multilayer networks

To provide usable services to the end user a telecommunications network consists of a number of layers, each providing features required to provide the final service. SDH is the transport layer that provides a manageable interface to the physical capacity. It can be seen as part of a simplified multilayer network as shown in Fig 1. End users provide demands that must be satisfied by the layers below, all of which impose limitations and restrictions, such as available resources, or technological limitations. For example, SDH circuits can be carried in wavelengths of which there are a finite number per fibre, and which can only traverse available installed fibre. The fibre topology is similarly dictated by duct availability. The layers are therefore closely coupled and highly dependent on each other.

## 3. Large-scale networks

Of the many layers in such a multi-layer network we will consider only two, together with an example network of each.




Fig 1 A simplified multilayer network.

### 3.1 The Internet

The largest example of an IP network is the Internet. It consists of a large number of interconnected IP networks that are not strictly planned and often considered near random in topology. Internally the sub-networks use interior routing protocols such as open shortest path first protocol (OSPF) or routing information protocol (RIP) to route along the shortest number of hops or to minimise other routing metrics. Internetwork routing is governed by protocols such as border gateway protocol (BGP) which implements a set of operator-defined policies to allow routing of traffic between autonomous system (AS) domains. The Internet has evolved and experienced huge growth but has no planned global structure or design and yet has been shown to exhibit a number of emergent topological traits. Faloutsos et al [1] examined a single instance of a router topology and a number of instances in time of the AS domain topology. They found that the following four power laws held true of all the topologies:

- Power-Law 1 (rank exponent)

The outdegree (connections from a node) was found to be proportional to the rank of a node, to the power of a constant, the rank being the position of the node in a table sorted (numerically decreasing) by the outdegree of the node:

$$d_v \propto r_v^R \quad \text{..... (1)}$$

- Power-Law 2 (outdegree exponent)

The frequency of an outdegree,  $f_d$  is proportional to the outdegree to the power of a constant,  $O$ :

$$f_d \propto d^O \quad \text{..... (2)}$$

- Power-Law (approximation) 3 (hop-plot exponent)

The total number of pairs of nodes, within  $h$  hops of each other, is proportional to the number of hops to the power of a constant. This is more of an approximation since it only holds for values of  $h$  that are much less than the network diameter:

$$P(h) \propto h^H, h \ll \delta \quad \text{..... (3)}$$

- Power-Law 4 (eigenvalue exponent)

The sorted eigenvalues (decreasing order),  $\lambda_i$  of the adjacency matrix (an  $N$  by  $N$  matrix (where  $N$  is the number of nodes) which is 1 when the two nodes are connected and 0 otherwise) are proportional to the index,  $i$ , into the list, to the power of a constant,  $\epsilon$ . The power law was shown to hold for only the top 20 or so eigenvalues:

$$\lambda_i \propto i^\epsilon \quad \text{..... (4)}$$

These very distinct traits are not planned and must exist as the result of individual design decisions by network operators or be the result of routing protocols. RIP is a distance vector routing protocol that attempts to optimise routes based on the number of hops, while OSPF is a link-state routing protocol that optimises for a specified metric (the default in Cisco routers for example is link bandwidth). They are both capable of reconfiguring the network in minutes, rather than hours or days.

### 3.2 SDH networks

In contrast to the dynamic and relatively unplanned nature of the Internet, SDH transport networks [2] are meticulously planned on a global (macroscopic) scale. They provide end-to-end circuits with strict guarantees on the delivered levels of service (such as available capacity, delay or jitter). To provide these guarantees the network must be dimensioned and planned in advance with the demands on the network being known (or carefully forecast) beforehand. During the planning stage explicit structures, such as rings, are often enforced in the network to provide for the easier provisioning of resilience, as they are a simple way to provide alternate paths through the network. Capacity is also often allocated to form a hierarchy, which allows for the easier management and better growth potential of the network. Any changes to the network are strictly planned and controlled and happen on a weekly to monthly scale.




Fig 2 The planned BT SDH network.

3.2.1 BT’s SDH network structure

These design principles can be seen in practise in the generalised diagram of the BT SDH network in Fig 2.

The network consists of approximately four tiers, the lower tiers being of a ring nature, multi-homed into upper tiers, while the top tier is a collection of highly meshed nodes. However, the greatest influence on the location and connectivity of these elements is still by equipment and fibre availability, user distribution and collocation of adjacent layer equipment.

3.2.2 SDH site topology

To investigate the deployed topology and the possibility of emergent properties a single instance of the configuration of BT’s narrowband SDH network was considered. The snapshot consisted of end-to-end circuits, the bandwidth allocated to these circuits and the geographic position of all the nodes traversed by the circuit. In this study we will be considering the connectivity of the sites, rather than individual rack equipment. In our network topology, a node is considered to be a site and two nodes are linked if there are one or more circuit hops passing between them. Since we are only considering the core network all customer nodes have been removed as well as the links to these nodes.

The network consisted of a few thousand nodes and had a links-to-nodes ratio of 1.77. A number of the plots in this study have been normalised to protect certain network

information — the power-law fit exponents or correlation factors are not influenced but the y-intercepts are.

Figure 3 shows a plot of the degrees (the number of links connected to a node) of each of the nodes against their rank when sorted in decreasing numerical value of degree. The SDH network exhibits a relatively good conformance to the power-law-fit with an  $R^2$  value of 0.929 ( $R^2$  is a standard statistical measure of deviation of a fit to data points with 1.0 being an exact fit). This is equivalent to Power-Law 1 in the Internet topologies. Degree rather than in-degree (number of incoming links) or out-degree (number of outgoing links) is being considered, as all circuits are symmetric.




Fig 3 Power-Law 1 — the degree rank exponent.

In Fig 4 we have the frequency distribution of the degrees of the nodes (Power-Law 2). We can see that again there is good conformance ( $R^2=0.932$ ) to the power-law-fit. This suggests that the network has developed a tendency to have a very large number of nodes with low connectivity and a small number of nodes with a very high connectivity.




Fig 4 Power-Law 2 — the frequency distribution of node degree.

In Fig 5 we have a plot of the number of pairs that are separated by the given number of hops (Power-Law 3). This law is only an approximation as it only holds when a node pair separation is 7 hops or less.




Fig 5 Power-Law 3 — the hop-plot approximate power-law (only holds when number of hops between nodes is 7 or less).

Power-Law 4 in Fig 6 is the eigenvalue power-law which is a plot of the eigenvalues of the adjacency matrix sorted in decreasing order plotted against the index into that list. Eigenvalues are a function of the network size, the cliques (sub-graphs) present in a topology, and their connectivity. The SDH data shows very good conformance to the power-law fit ( $R^2=0.995$ ), even for the thirty highest eigenvalues.

We can see that, while created very differently, the Internet and this SDH network both exhibit the same emergent topological traits and follow the same power-laws, albeit with different exponent values. The SDH network does, however, follow further power-laws.




Fig 6 Power-Law 4 — the descending sorted list of real eigenvalues of the adjacency matrix.

In an effort to find the effect of imposed structure it was found that the clustering co-efficients [3] of the nodes formed a power-law with their rank in a sorted list. The clustering co-efficient of a node is defined as the ratio of the total number of links between all of the nodes in the neighbourhood of a node (the node and all of its immediate neighbours) and the total number of possible links in that neighbourhood:

$$c_v = \frac{l_{neigh}}{0.5d_v(d_v + 1)} \quad \dots (5)$$

where  $c_v$  is the clustering co-efficient of node  $v$ ,  $l_{neigh}$  is the number of links between all nodes in the neighbourhood and  $d_v$  is the number of nodes directly connected to the central node,  $v$ .

The plot of co-efficient against rank can be seen in Fig 7. It has a number of interesting features, including the upturned tail on the left and the plateaus in the right half. The tail could be the result of the removal of customer nodes and links, which would cause an increase in the clustering co-efficient of edge nodes since they were previously hub nodes. The rightmost plateau, which is at a clustering co-efficient of 1.0, is probably the tier 1 nodes which are highly meshed. The much wider plateau of 0.66 is probably caused by homogeneous rings in the lower tiers as there so many of them.

The fifth power-law can therefore be defined as:

- Power-Law 5 (clustering co-efficient rank exponent)

The clustering co-efficient,  $c_v$ , was found to be proportional to the rank of the co-efficient, to the power of a constant,  $C$ :

$$c_v \propto r_c^C \quad \dots (6)$$

We have now seen how a globally planned network exhibits a range of topological traits. To further understand




Fig 7 Power-Law 5 — the rank plot of the clustering co-efficients of all the nodes.

the sources of these traits and the extent to which they are present we will consider the geographic distribution of the SDH sites.

### 3.2.3 Network geography

One of the aims of network planning is cost minimisation and as a result network access points are often placed as near to the centre of clusters of customers as possible. This minimises costs associated with customer access systems and therefore causes the geographic site distribution to follow population distribution. It has been previously shown that the customers, and therefore SDH nodes, follow a fractal distribution [4], but we must examine the effect that the distance between SDH sites may have on their connectivity. In Fig 8 we find the distance between every node pair plotted against the probability of a link existing between nodes over that range. The distances between nodes are grouped in bins of 500 distance units.




Fig 8 The probability that a pair of nodes are connected versus the geographic distance between the two nodes (distances have been put in 500 distance unit bins).

We can see that there is a sharp drop in linking probability as the distance between nodes increases. Since there is a clear relationship between link probability and distance the connectivity must be a function of node distribution and therefore physical layer cost functions.

Now we will consider how prevalent the traits are throughout the network geography and the network hierarchy. As we move from town, to city, then national scale we follow the tiers up the hierarchy of Fig 2. To examine the existence of these traits throughout the tiers we can divide the country by a grid of fixed sized squares. The width of the square will be denoted by  $M$ , which is in the distance units used before. As  $M$  increases, more and more nodes are grouped within these squares. The connectivity of the squares then follows the connectivity of the nodes enclosed within the squares. As  $M$  increases each square will encapsulate towns, then cities, counties and finally the whole country will be contained within a single square.

In Fig 9 we see the effect of increasing  $M$  on the exponent of the degree distribution (Power-Law 2). The right axis shows the  $R^2$  value of the power-law fit. It is not until  $M$  reaches 4000 distance units that the conformance drops below 0.90. At this scale we are examining entire towns, but large cities could still span multiple grid squares.




Fig 9 The Power-Law 2 exponent for increasing grid size,  $M$ , as well as the conformance to the power-law-fit,  $R^2$ , on the right axis.

So even the connectivity of entire towns, containing many SDH sites follows Power-Law 2. The extent of this can be further seen in Fig 10 where  $R^2$  is plotted against  $M$ ,




Fig 10 The conformance to Power-Law 2, for increasing  $M$ , with the fraction of total number of nodes formed by the grouping of the grid, on the right axis.

but with the fraction of total nodes in the grid on the right axis. The conformance does not drop below 0.90 until the topology contains 10% of the original number of nodes.

The grouping of nodes into squares at different scales (values of  $M$ ) is equivalent to the box-counting [5] method for finding the fractal dimension of the geographic node distribution. The clear slope between  $M=1000$  and  $M=10\,000$  demonstrates that site distribution is scale invariant and the gradient of this slope, and therefore the fractal dimension, is 1.46.

### 3.2.4 Bandwidth distribution

Another aspect of the network to be influenced by network geography and topology is bandwidth distribution. User distribution, and therefore demand end-points, together with routing algorithms and the available physical capacity, all dictate the distribution of allocated bandwidth within the network. In Fig 11 we can see this distribution, where total allocated capacity between nodes, rather than available physical capacity is plotted.




Fig 11 Allocated bandwidth distribution between sites (normalised bandwidths have been put in fixed-size bins).

This distribution follows a power-law with a weaker conformance of 0.873 and has a slight tendency to arch on the left half of the plot.

We have now seen how a globally planned network, designed by well-specified algorithms and under the influence of undefined forces has formed a number of macroscopic emergent traits. The traits do not appear explicitly in network design algorithms and therefore must be the response of the algorithm to the input data set.

## 4. Where do these emergent properties originate?

The emergent traits of the SDH layer are unplanned and unexpected. The source of the traits must therefore lie in the formation and evolution of the network and unspecified local effects on the growth. SDH networks do not use dynamic routing and when changes do happen they occur on weekly or monthly scales and are often

incremental rather than entire reconfigurations. The major influences on these changes are fourfold.

- Technology

There are very large restrictions on the planner, from equipment limitations, to capacity limits, capacity granularity, physical connectivity and legacy network demands. The BT SDH network also has to cater for the large amount of legacy traffic that was originally plesiochronous digital hierarchy (PDH) based. The original planning of this traffic was decided by PDH planning techniques that had even more technological restrictions, such as those imposed by the 'multiplexer mountain'. In contrast to this, IP networks are much more dynamic and far less inhibited, with links capable of crossing the entire country through a single SDH or ATM circuit and there being no capacity granularity issues. IP dynamic routing protocols, which can reconfigure the network in seconds or minutes, often have metrics based on either link capacity (default OSPF metric behaviour in Cisco routers) or hop count (RIP) to destination. SDH planners consider a combination of the two. These technologies, routing protocols or network architectures, do not, however, impose explicit power-law traits.

- Network architecture

When adding new links to an existing SDH network the planned structure must be followed. It can be shown, however, that simple models of an SDH network such as a three-tier random collection of rings of nodes does not exhibit such laws, neither does a three-tier collection of networks which are internally randomly connected (a simple model of the Internet). Aside from the basic connectivity elements (rings, meshes) the two network types differ further — in organisation. The Internet consists of stub networks that contain end users, and transit networks that carry the traffic between the stub networks [6]. In SDH networks on the other hand 97% of the sites are connected directly to an end customer; there is little distinction between edge nodes and core nodes. It is conceivable therefore that the topology of the SDH network is much more closely coupled to the demand of the layers above than an IP network would be.

- Adjacent layer cost functions

The design of each layer must be accommodated by the layer below and as such the design will be influenced by cost functions imposed by that layer. Due to its proximity to the physical layer SDH is much more tightly coupled to geography. IP on the other hand makes demands of the transport layer and is less inhibited by the physical layer: a single hop in an IP network could be carried by a single SDH circuit across the length of the entire country. IP does not

necessarily need to be carried over SDH either, and could use a range of layer 2 technologies, such as frame relay, FDDI or Ethernet, all of which have their own limits.

It can be shown though that cost-function-based connectivity models for geographic distribution do not necessarily produce power-law compliant topologies. Based on uniformly random node distribution, or even heavy-tailed node distributions and a connection cost function such as that proposed by Waxman [7], the traits do not exist.

- Demand growth characteristics

The specifics of the demands being added could result in these traits. Different types of customers will exhibit different types of growth. With the growth of Internet usage at home, residential customers may require more capacity to reach Internet service provider (ISP) points of presence (PoPs), or with the growth of virtual private network (VPN) services there will be increased demand between business sites. If the demand problem is simplified to assume that the new nodes will have an affinity to the larger of the existing nodes, then we could model such growth using a prototype by Albert and Barabási [8]. Their models concentrated on preferential connectivity according to existing connectivity; new nodes would connect to existing nodes favouring the more connected nodes. The models do in fact produce topologies that follow power-laws. The exponents of the fits are, however, different to the exponents measured for the Internet and BT SDH network, but, more importantly, they are too simple as they totally disregard cost functions such as those associated with geographic distribution and technological restrictions. Additionally, for SDH networks, they do not consider the explicit network design.

Of these possible influences only the growth model by Albert and Barabási could produce the power-law traits. The effect of the other influences is undeniable however: The network architecture does exist, as do cost functions (Fig 8) and so do technological limitations. For further models producing power-law topologies, see Mitzenmacher [9], and on the effect of structure on topologies, see Tangmunarunkit et al [10].

#### 4.1 Self-organising criticality and multilayer feedback

When discovering the emergent traits of the Internet, Faloutsos et al suggested [1] that the source could be a self-organising system that would reconfigure the topology to cater for increased network demand. This could imply that core nodes with their higher capacity links will also have higher connectivity. To investigate this,

Fig 12 has a plot (for the SDH network) of the bandwidth of a link against the degree of the less connected of the two end-points of the link. There is a very concentrated cluster in the lower left quadrant of the graph and a lack of high bandwidth with low degree links, the correlation coefficient of the points is 0.611.




Fig 12 The lower of the two degrees of the link end-points versus the allocated bandwidth of the link.

This would suggest that, if there is a relationship between node size (capacity) and its connectivity, it is not a simple one. This being the SDH network, the lack of an obvious relationship could be caused by imposed structure or the lack of a clear distinction between edge and core nodes. The lack of a clear relationship does not, however, discount the theory; it can be shown that the addition of core links to sustain increased demand from a growing network also creates power-law adherent networks [11].

The problem again with this growth model is the lack of physical layer cost functions and network structure. In Spencer and Sacks [11], the use of physical layer cost functions immediately made the network deviate from the power-law traits.

The source of the traits is the subject of on-going research and the contributions of all of the above sources must be considered. The universality of the traits, being in different layers and at various levels (i.e. it exists in both Internet router and AS domain topologies) of each network suggests the source may lie in a more inherent process.

It has been suggested that as part of the self-organising system the need to minimise resource usage could be responsible [12]. That would certainly be present in both SDH networks, as part of the planning algorithm, and the Internet, as part of metric-based routing protocols. If the dynamics of layer growth and its effects on adjacent layers are to be properly understood, then better models for these traits must be devised.

## 5. Discussion and conclusions

We have presented a number of unplanned traits emerging from a strictly planned and highly structured SDH network. The traits are present at the site topology level and continue up to the connectivity of towns and cities. While this was an examination of only a single instance of an SDH network, the mere possibility has very wide reaching implications.

In traffic forecasting and scalability analysis, knowledge of the macroscopic design can be used to find bottle-necks and under-utilisation. These traits also give a better picture of the demand an SDH network can place on the physical layer and WDM planning can now benefit from more realistic demand models, as would simulation of the SDH network itself [13]. If the sources of these traits were properly captured then it would be possible to predict the effect of physical layer changes on the layers above and develop more robust and scalable networks.

The actual traits, and not just modelling based on them, have wide ranging implications too. As there are no size-dependent characteristics in the plots, the network is considered as not having obvious limiting factors to its scalability. The precise source of the traits is elusive and is conceivably part of the large-scale problem of resource minimisation within multilayer networks and the feedback between the constituent layers. Layer limitations, design principles, legacy network contributions and population interaction all play a part in this self-organising feedback system and their contributions are hard to evaluate.

## References

- 1 Faloutsos M, Faloutsos P and Faloutsos C: 'On power-law relationships of the Internet topology', in Proceedings of SIGCOMM: IEEE Conference on Communications (1999).
- 2 Brown R S, Rowland D, Vinall A and O'Neill A: 'Broadband transport — the synchronous digital hierarchy', *BT Technol J*, **16**, No 1, pp 148—158 (January 1998).
- 3 Watts D J and Strogatz S H: 'Collective dynamics of 'small-world' networks', *Nature*, **393**, pp 440—442 (June 1998).
- 4 Appleby S: 'Fractal telecommunication networks', *BT Technol J*, **12**, No 2, pp 19—29 (April 1994).
- 5 Peitgen H O, Jürgens H and Saupe D: 'Length, area and dimension: Measuring complexity and scaling properties', chapter in: 'Chaos and Fractals: New Frontiers of Science', New York, Springer-Verlag (1992).
- 6 Calvert K, Doar M and Zegura E: 'Modelling Internet topology', *IEEE Communications Magazine* (June 1997).
- 7 Waxman B M: 'Routing of multipoint connections', *IEEE Journal on Selected Areas in Communications*, **6**, No 9, pp 1617—1622 (1988).
- 8 Albert R and Barabási A L: 'Topology of evolving networks: local events and universality', *Phys Rev Lett*, **85**, No 24, pp 5234—5237.
- 9 Mitzenmacher M: 'A brief history of generative models for power law and lognormal distributions', 39th Annual Allerton Conference on Communication, Control, and Computing (2001).
- 10 Tangmunarunkit H et al: 'Network topology generators: degree-based vs structural', *SIGCOMM*, IEEE Conference on Communication (2001).
- 11 Spencer J and Sacks L: 'Modelling IP network topologies by emulating network development processes', *IEEE International Conference on Software, Telecommunications and Computer Networks* (2002).
- 12 Fabrikant A et al: 'Heuristically optimized trade-offs: a new paradigm for power laws in the Internet', *International Colloquium on Automata, Languages, and Programming (ICALP2002)* (July 2002).
- 13 Anagnostakis K, Greenwald M and Ryger R: 'On the sensitivity of network simulation to topology', *Proceedings of the 10th IEEE/ACM Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS)* (2002).



Jason Spencer obtained his BEng in Electronic Engineering and MSc in Telecommunications in 1997 and 1998 from University College, London.

He is currently completing his PhD at UCL on the interactions between network layers and their effect on network design.

His research interests include network planning and management, de-centralised network control, next generation high-speed metropolitan networks, complex systems and large-scale system design.



Dave Johnson joined BT in 1986 from Lancaster University with a first class honours degree in Physical Electronic Engineering.

He has worked on systems reliability, self-organisation in networks, and the application of novel network design and planning techniques. His work on resilience strategy led to the deployment of SDH shared protection rings in BT's network in 1999.

He is currently Core Transport Architecture and Design manager for BTextact Technologies. He leads a team of engineers developing pragmatic solutions to transport network problems. This includes the design of new networks, cost and performance modelling, supplier selection, and the development of transmission strategy.

## Emergent properties of the BT SDH network



Andrew Hastie joined BT's Network Architectures and Design group at Adastral Park in 1999, after being awarded an MEng in Electronic and Electrical Engineering with Business Studies by Strathclyde University.

On joining BT, his work involved gathering information on BT's SDH network to make recommendations about circuit routing rules and their impact on availability. He was also involved in the design of BT's core broadband STM-64 network.

In more recent times he has been involved in defining and measuring the utilisation of the SDH network, while gathering network data to define the efficiency of circuit routings, making the identification of mis-routed circuits easier.



Lionel Sacks became a lecturer in the Department of Electrical and Electronic Engineering at UCL in 1997. His research is on telecommunications systems, with a particular focus towards high integrity and high availability. His research group (ACSE) is involved in work exploring autonomous self-organising systems as applied to areas such as network traffic, topologies management, service engineering, sensor networks, knowledge management and work distribution for computationally supported services (grid and active networks). Prior to working in engineering, he trained and

worked in high energy physics having undertaken his PhD on the experiment CERN/NA1 measuring charmed particle lifetimes and spent 7 years at Liverpool University and CERN on CERN/PS195 measuring CP violation.