Requirements for the Next Generation of Building Networks

Peter Palensky
Institute of Computer Technology, Vienna University of Technology
A-1040 Vienna, Austria

ABSTRACT

Today's building automation and control systems (BACS) are based on networks that can cope with

- very large numbers of nodes,
- flexible and scalable topologies, and
- sophisticated network management.

The increased need for connectivity, communication and remote access results in plenty of other network systems, found in buildings like

- telecommunication,
- office automation, or
- safety and security equipment,

to name only three of them. The widely desired convergence of these networks leads to substantial problems. Issues like real-time capabilities, costs per node, management or the topology are addressed in this paper.

It gives a wish list of features, taken from existing network technologies from industrial automation and other domains, that should be combined, in order to create a network that is capable of satisfying the needs of future buildings.

The goal is to have one consistent and vertically integrated technology for all application. This compromise must offer a level of flexibility, that can very easily lead to a high system complexity. Therefore scalability and simplicity are the commandments that should guide the design of the next generation of building networks.

Keywords

Building automation, networks, protocols, convergence, control networks.

1. INTRODUCTION

Modern homes and buildings are getting more and more "networked". This is true for private homes as well as for large office buildings. The driving forces behind this trend are amongst other things the need for

- more and better information, for
- more mobility, for
- increased "performance" of buildings, and for
- more flexibility.

Large office buildings are often equipped with a building automation and control system (BACS) that manage heating ventilation and air conditioning (HVAC), lighting, sun blinds and other things. Even if the initial costs for the equipment is sometimes triple the costs of traditional electrical wiring, it amortizes when the the building is in operation. It can be run more efficiently since energy costs can be reduced dramatically with the right energy management system.

Energy management is typically divided into "energy efficiency" and "demand response", where the former means replacing old and inefficient equipment, insulation and the like, and the latter is the reaction of the demand side (the customer) to changes in the energy price. Generally this demand response means "consume energy when it is cheap and try to reduce consumption otherwise". This strategy can be automated by means of control systems that schedule the operation of equipment. This is commonly known as "demand side management" or DSM.

DSM tries to avoid peak loads, shift consumption to other times, etc. without without influencing the performance or output of the overall customer process. The basis for this are "virtual energy storages". Any thermal (heating, cooking, cooling, freezing, etc.) equipment can store energy in form of heat. So if the interval between 11:30 am and 01:30 pm is too expensive, the system can try to heat the building before 11:30 that much, so that it can stay unheated during this high-price time. Naturally, the system has to respect many side-effects in order not to disturb the actual "purpose" of the building, like being a comfortable office. Beside these thermal storages, there are further virtual energy storages like air conditioning where the percentage of carbon dioxide in the air is the storage or organizational ones, where some schedule can be re-organized.

A BACS is the only reasonable technology for implementing DSM in a building. A control system – if it is a centralized or a decentralized one – needs remote access to all relevant energy consuming equipment in the entire building, to energy counters, and it needs on-line pricing information for its algorithm. A BACS that was originally only intended for operating the blinds or lights can easily be used for that.

This leads us to the main advantage of open and standardized building automation and control systems: "add-on-services". One and the same infrastructure can be used for a variety of applications. A BACS might be too expensive for simply replacing an ordinary heating control system. Using the same wires and the same management software to make an on-line analysis of the situation of the building (energy consumption, water leakage, status of windows, etc.) is often the last missing reason to convince someone of the advantages of such systems.

Another reason for using BACS is the increase of flexibility. Buildings are changed during their life cycle, they are extended, rooms are added, or merged, walls are moved, etc. Traditional electrical installations would have to be completely removed and redone while BACS simply need changes in their software configuration: it is easy to tell a BACS switch to operate other lamps than the day before.

BACS, however, are not the only networks that can be found in buildings. The requirements of a fire or burglar alarm system might lead to expensive network technology that can not be afforded for every light switch. A recent trend are multimedia networks used for sound and video streaming. They are used in private homes for entertainment as well as in office buildings for video conferencing or video surveillance and conferencing. The goal is to have any multimedia content available wherever you are in the building. Such networks have - compared to BACS - a tremendous need for bandwidth and other aspects of Quality of Service (QoS). Applications like "Video on Demand" or "Voice over IP" (VoIP) therefore lead to broadband networks like IEEE 1394 (also known as "Firewire") or Gigabit Ethernet.

A further and similar type of network that can be found in buildings is the office network typically based on a local area network (LAN). LANs connect workstations with each other, with the Internet, with corporate servers and with other IT equipment like printers, scanners, etc. Recently, LANs are more and more converging with multimedia networks, having the Internet Protocol (IP) and Ethernet as a common basis.

The fourth and last network in buildings is the telecommunications (telecom) system that connects telephones and intercom systems.

All these networks co-exist, have their own technologies and their own and separated market.

2. PROPERTIES OF NETWORKS IN BUILDINGS

The four types of networks,

- BACS,
- Multimedia networks,
- Office networks, and
- Telecommunication networks

differ in many ways which is the reason why they (still) exist side by side without replacing each other.

Physical channel

Multimedia networks and telecom networks need guaranteed bandwidths since they provide interactive and session-oriented communication. The majority of BACS and office applications can get along with stochastic delays, as long as they do not exceed some critical (statistical) level.

Topology and number of nodes

BACS sometimes have tens of thousands of nodes, while - even in large buildings - one multimedia node per room may suffice. The physical sizes of the networks are meanwhile pretty much the same, only BACS are more fine grained, since they go down to the last sensors.

Network management

Multimedia networks - and especially personal area multimedia networks - face the situation that their nodes "come and go" when for instance a digital camera is plugged in and out (the same holds for wireless access points). They are specialized in plug-and-work and automatic reconfiguration, while BACS are typically "managed": someone has to register and administer nodes in some sort of database. Management of such large networks is not an easy task and demands sophisticated network management tools.

Security and safety

Network security means that networks offers the right means against the following threats [1]:

- Confidentiality - unauthorized network members should not be able to listen to or read out data, typically achieved by encryption.
- Authentication - the communication peers should have mutual certainty about their identity. A digital signature can be used for this.
- Integrity - alterations of transmitted data should be identifiable. This can be achieved by using an electronic envelope.

While office networks nowadays show the first steps towards this, the others usually have no security means at all (this is especially true for BACS [2]).

Safety of networks leads to galvanically insulated components, fail safe states for all nodes, robust media, redundant architecture, etc. and a methodology to quantify this safety for insurances. Safety-relevant applications typically use expensive networks that guarantee their function up to some certain level. In buildings, safety is getting more and more relevant and companies are currently forced to use technologies that were definitely not designed for buildings, like ASI, Profibus or CAN, to implement safe BACS.

Quality of Service

Certain applications make high demands on QoS. Control algorithms for instance rely on a guaranteed sample rate or delay if the parts of the control loop are interconnected via a BACS network. Telecom networks on the other hand might desire billing for certain services with different quality and availability.

3. A WISH LIST

The technology and the market for these networks is now mature and established and it is time to think about the next step: the convergence of these networks.

The wish list for a future building network technology is the following:

Multi-layer Interoperability

A digital TV in the living room set should "understand" a temperature sensor in the cellar. All nodes should be able to provide or consume services and data to or from other nodes. The often mentioned "IP-down-to-the-sensor" phrase is one expression of the need for vertical integration and transparent network services.

Scalable and flexible Topology

The network should have virtually no physical limits like node address range, length of the wires or a restricted topology. This calls for sophisticated routing, electrical characteristics and network management. Adding a probably large new segment to the network must not run the system into troubles. Additionally it should be possible to have cheap and simple nodes side-by-side with high-performance nodes that may cost more.

A combination of plug-and-participate and network management

It must be possible to add and remove individual nodes without consulting a system administrator. If the application, that node takes part in, is easy enough, it should even offer plug-and-work. This plug-and-something must seamlessly integrate in and cooperate with traditional network management and the corresponding tools.

Usage of off-the-shelf technology

Ideally the new network needs no development at all but uses and combines existing technology that is proved, accepted and affordable.

Dramatic reduction of hardware costs

The nodes of the network must be as cheap as possible. It must be possible to build sensors for less than one dollar without losing functionality or quality, since the initial costs are still one of the main obstacles for building networks.

Integration with existing networks

An existing Ethernet-based LAN can probably be part of the new network. This re-use of existing infrastructure greatly reduces costs and increases acceptance.

Robustness

The network should be robust, available and reliable. Self-healing mechanisms and fault-tolerant and "gracefully degradable" behavior is very much wanted. The network topology should therefore allow a mesh. Having alternative paths for packets can also be used for "reverse multiplexing" where parallel channels are merged to one high-performance channel which might be especially useful for multimedia content [3].

The least common multiple of existing features

The special features of the individual networks now should not be lost. The new network must be able to satisfy real-time requirements of a control system as well as high scalability. If the features contradict each other, a smart compromise must be found, that works for both applications.

4. THE NEXT GENERATION OF BUILDING NETWORKS

Fulfilling all the above mentioned requirements is the objective of the next generation of building networks. In order to achieve real Interoperability, the nodes must settle on the same data types, formats and ideally on the same protocol stack. Gateways between two networks "lose" or alter the quality of information when forwarding it from one side to the other. Having the same protocol for all applications (office interconnection, BACS, multimedia, etc.) would release us from gateways that cause immense maintenance costs and network complexity. The network technology that we should aim for should be able to interconnect any node in a building with any other node. This network needs:

A scalable protocol stack

Scalability is necessary in all protocol layers from the media arbitration up to Interoperability definitions above layer 7 according to the ISO/OSI 7-layer reference model.

With a scalable protocol stack it will be possible to send packets from a couple of bytes up to large protocol data units (PDUs) that are used in multimedia applications. Nodes with extreme low costs can not contain a sophisticated microprocessor that operates a complicated protocol stack or a buffer that is able to store tens of kilobytes.

Therefore a variable-length frame (layer-2-PDU) is needed, that goes much further than it is done in Ethernet or frame relay where the payload is the only variable part [4].

Such a frame must start with a header where already a few bits tell its "class" or size. Simple light switches might therefore only send 3-byte-frames while sharing share the channel with nodes that exchange kilobyte-frames. Theses simple nodes might not need any buffers and can be implemented by using an extremely low-cost programmable chip like a gate array logic (GAL).

The frames can for example be divided into three classes like

- "A": 3 bytes frame, similar to or even more primitive than I2C or SPI protocol,
- "B": 10-64 byte frames, similar to existing BACS like LonWorks or KNX, can be handled by controllers in the 8051 or PIC category, and
- "C": 64-1024 bytes frames, the size of Ethernet and BACnet.

A device of class "B" can understand "A" and "B" frames but needs a proxy for talking to "C" nodes. (see section "Scalable topology and routers" in this chapter). "C" frames on the other hand will be ignored by "A" and "B" nodes.

All these frames have to contain delimiters, a header, address information, checksums and the payload itself. The 24 bits of the 3-byte class "A" frame therefore might only be 1 start-bit, 2 bits header (to distinguish it from class "B" or higher), 1 parity bit, 8 bits for the node address, 3 bits for the register address (within one node), 8 bits payload, and 1 stop-bit.

An example for scalable Class "B" addressing can seen in P-Net (an automation network standardized in EN50170, mainly used in food processing): the length of the address is variable and the number of network segments is virtually unlimited. The PDUs move - guided by an address in their header - worm-like from router to router until they reach their destination [5]. The address information in a class "C" frame can use standard methods, known from Ethernet/IP.

Layer 3-7 should allow implementing only parts of the protocol stack. Like class "A" nodes do not have layers 3-7, there should also be the possibility to implement only parts of the stack, in order to decrease the memory-footprint of the nodes. If for instance the protocol stack defines various security services or a session service, it should not be obligatory to implement this in every node.

The same holds for network management services. Some nodes might offer remote-diagnosis, remote application management and similar management services, but others might not. Imagine two temperature sensors, where one sends its value every 5 minutes while the heartbeat-rate of the other one can be configured via network management functions.

In order to know what a node is capable of, it might become necessary to introduce sub-classes where for instance "C5" contains all "C4"-services plus a couple of additional ones.

Above layer 7, there must be a scalable version of Interoperability rules. Some networks, LonWorks for instance, store device documentation, node-identification, data types, and other descriptive date almost entirely on the nodes. In this distributed way, every node has all necessary information on-board. Other systems rely on an external database that tells how what nodes can interact.

We can not expect class "A" nodes to contain interoperability and self-documentation information like an SNMP-able IP node [6]. Other nodes, preferably the routers, have to support these primitive members of the network to get a quality of interoperability that one is used to in industry, where guidelines like CANopen (EN 50325-4, [7]) or LonMark [8] make life easier.

Hybrid arbitration

Unlike ring-based systems [9] or point-to-point networks, the new network system should be able to manage a shared communication media. Two entirely different strategies to access this media are event-triggered and time triggered protocols. The former ones simply try it, while the latter ones usually have certain time slots assigned. Event triggered access can cause collisions when more than one node decides to access the bus. There are various methods in use to detect and even avoid collisions.

Such CSMA (carrier sense multiple access) networks are easy to extend. Just add a node and it works - until collisions make the channel collapse. Plain Ethernet's CSMA/CD (CSMA & collision detection, [10]) for instance has no means to avoid this, predictive p-persistent CSMA of the LonTalk protocol is one step better, it uses slots with priorities to arbitrate and adapts itself to the expected network load [11].

A simple time-triggered protocol would divide the bandwidth into n time slots where each node has exclusive access to its slot. The problems occur, when the n+1th node is added to the network. Such slot-based arbitration might be real-time able, easy to implement and reliable but wastes bandwidth and does not scale well. Examples for this can be found in industrial automation ([12], [13]). Also token passing methods are widely used [14]. While systems like the one described in [15] only respect periodic real time requirements (while still having collisions!) the new system has to take care about sporadic and spontaneous hard-realtime bandwidth needs as well, probably only in a statistical way [16].

An interesting combination is described in [17]: hard-, soft- and non-realtime requirements are combined in one (drive-by-wire) application. Static and dynamic scheduling mechanisms guarantee hard real time requirements while still not waisting bandwidth.

The next generation of building networks should try to combine the advantages of both worlds. But unlike HYMAP [18], where the arbitration method changes depending on the network load, the new protocol should permanently use event- and time-triggered methods. Parts of the bandwidth should be assignable via slots, while the rest should be used for CSMA. This strategy is already successfully used in IEEE 1394 networks or ATM. Also Flexray, a TDMA (Time Division Multiple Access) based network, divides its bandwidth into a static segment and a dynamic segment to have the advantages of both worlds.

The desired protocol should again go some steps further. The slots should be "booked ahead", the bandwidth should be allocated dynamically and anticipatory. If for instance a streaming service needs a certain bandwidth for a certain time, it should only reserve slots for the time it needs it. The nodes should be able to reserver this bandwidth "on-demand" similar to an IEEE 802.16 MAC (media access) sublayer [19].

The CSMA part should respect priority messages for critical applications. Non-priority messages should have a Gaussian and probability for bus access and should completely avoid collisions, similar to CAN (control area network) where recessive bits determine who gets the bus [20]. Also fairness is an important issue, which might be done by using a credit-based bandwidth-on-demand method similar to [21].

In this way, a sophisticated CSMA/CA (CSMA collision avoidance) coexists with dynamically reserved slots. The main challenge is that simple nodes (the "GAL" class) should still be able to take part in this arbitration procedure. You can not expect these nodes to participate in complicated voting mechanisms for the assignment of slots. This assignment of bandwidth must be done by more powerful nodes, supported by a network management system.

Scalable Topology and routers

The network must offer the possibilities of having individual segments interconnected via routers that forward messages and translate the possibly different bitrates (buffering, intermediate acknowledgment messages, keep-alive packets, etc.) and different media (RS485, Ethernet, etc.). Imagine a segment in a single room that is mainly used by primitive class "A" nodes like light switches, a temperature controller, a constant light controller, etc. The controllers and their dedicated sensors and actuators needs to reserve bandwidth for their control application. Luckily this need for bandwidth is only within this single room: the sensors, the controller and its actuators are all in one and the same room. Consequently there is no need to reserve the bandwidth in the entire building but only in this room that is coupled via a router. Bandwidth reservations that span routers are not that easy. They demand their reservation also in the transit net (the path between the nodes).

Simple class "A" PDUs do not know about routers, or layers 3 or 4 in the ISO/OSI reference model. They can only address nodes in their segment. A router therefore has to mirror distant nodes to the internal segment. In this way, a temperature sensor in the cellar can send its value every minute to a class "C" PC in the living room, as long as the PC appears with a local class "A" address on the internal port if the router: the router acts as two-way proxy that reflects the relevant part of the rest of the world.

Despite its partial simplicity, the network should be a peer-to-peer network that avoids master/slave relations (like done in LIN - local interconnect network -, a very simple and low-cost network for automotive applications [13]).

Additionally, the network should be capable of maintaining multiple self-healing paths from A to B. If a segment breaks, the really important packets should still be able to use alternative paths, even if a Gigabit Ethernet channel is then bridged via a 78kbit/s RS485. The respective routing algorithms should be able to follow various rules like energy-aware routing, least-cost routing, load balancing and the like.

Network management

Network management is a crucial point, especially because in buildings we have to expect extremely high numbers of nodes. What is needed is a consistent way of using distribute (i. e. on-chip) network management features and centralized ones. Additionally the system should utilize new possibilities.

If the channel is a wireless one, the nodes can usually determine the distance to other nodes by using the strength of the signals. This can be used to determine their position via triangulation. Position information is a very useful thing if 2 new and identical nodes (e. g. light switches) are discovered in one segment and no-one knows which one is which one. Having the physical 3-D location can be used to assign the right configuration data like which lamps to switch (by fining the node in a CAD-schematic by even guessing that it will switch the lamps in the same room). Non-wireless nodes can be equipped with RF-IDs in order to locate them.

Defective nodes should be automatically registered in a maintenance protocol. If a node is replaced, the new node must be integrated as automatically as possible. This implies that the logical address of the new node should be the address of the replaced one. The assignment of logical addresses (e. g. IP addresses, subnet/node address, etc.) should support plug-and-participate as well as manually managed assignment.

Flexible Physical Channels

The new network must support a variety of physical channels, some specialized for rough environments (LonWorks' FTT10 for instance is galvanically insulated and extremely robust regarding EMC), others for high bandwidths (fiber optics as a backbone).

Ethernet and IP as omnipresent technologies in buildings should coexist with (if not even be the basis for) the next generation of building networks. Industry already uses Ethernet to encapsulate existing fieldbus messages. Ethernet/IP uses the Control & Information Protocol (CIP) that runs on top of TCP/IP and UDP/IP [22]. Similar software solutions are done in ProfiNet or Modbus-TCP.

Real time extensions for Ethernet like "Ethernet Powerlink" [23] offer an isochronous and an asynchronous channel, similar to IEEE 1394. Another example is given in [24], where - again - a software layer implements a token- or slot-based mechanism on top of the Ethernet's natural CSMA.

The main obstacle is that the sophisticated arbitration of the new nodes is spoiled when traditional Ethernet nodes join the segment. This problem is known and usually solved with a switch. A "layer 3 switch" can for instance be used to guarantee QoS for VoIP nodes. The IPv4 header contains a ToS (Type of Service) field whose DSCP (Differentiated Services Code Point) bits can carry a priority that is respected by the switch. In this way, a switch can separate channels that would otherwise interfere with each other. IPv6 offers even more sophisticated QoS possibilities.

An exotic strategy to achieve low-cost nodes is to allow multiple bitrates on one physical channel. Very cheap RS485-nodes might not be able to talk faster than 9.6 kbit/s. It has to be guaranteed that all nodes can detect a "busy" channel (carrier sense) whatever bitrate is used, in order to use one consistent arbitration strategy.

All possible channels, however, must support a common set of basic services (like broadcast messages) independent of their technology.

5. CONCLUSION

The convergence of networks in buildings is not easy, otherwise it would already have happened. It is necessary to take networks from different domains (industrial automation, telecommunication, etc.) as examples on how various requirements can be met. Combining the various features under one umbrella is of course a challenge.

The key to success will be simplicity and reduced costs, two attributes that the new network must fulfill. Having scalability as a rule for all design questions helps reducing costs, since the entire network infrastructure can be tailored to its needs without violating the overall concept. Scalability is also the path also to an as-low-as possible system complexity.

Further publications will report on the first reference implementations, test results and the used technology.

6. REFERENCES

[1] P. Palensky, "Smart Card Security for Field Area Networks", **Proceedings of the IEEE-Siberian Conference on Communications and Control SIBCON-2003**, Tomsk, 1.-2.10.2003

[2] C. Schwaiger, A. Treytl, "Smart Card Based Security for Fieldbus Systems"; **9th IEEE Conference on Emerging Technologies and Factory Automation**, 2003, pp. 398 – 406.

[3] P.H. Fredette, "The past, present, and future of inverse multiplexing", **IEEE Communications Magazine**, Volume: 32 , Issue: 4 , April 1994 pp. 42 – 46

[4] Uyless Black, **Frame relay networks : specifications and implementations**, McGraw-Hill, ISBN 0-07-005558-0, 1994

[5] **The P-NET Standard**, International P-NET User Organisation ApS., 1994

[6] W. Stallings, **SNMP, SNMPv2, SNMPv3, and RMON 1 and 2**, Addison Wesley, 3rd ed., 1999.

[7] M. Farsi, K. Ratcliff, "Controlling with CANopen", **IEE Review**, Volume: 44 , Issue: 5 , 17 Sept. 1998 pp. 229 - 231

[8] **LonMark Application Layer Interoperability Guidelines V3.3**, LonMark Interoperability Association, USA, 2002

[9] A. Bhargava, J.F. Kurose, D. Towsley, "A hybrid media access protocol for high-speed ring networks", **IEEE Journal on Selected Areas in Communications**, Volume: 6 , Issue: 6 , July 1988, pp. 924 – 933

[10] "Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", **IEEE Standard 802.3**, IEEE Inc., New York, USA, 2002

[11] Loy, D., Dietrich, D. and. Schweinzer, H.-J (Eds.), **Open Control Networks, LonWorks/EIA 709 Technology**, Kluwer Academic Publishers, 2001

[12] E. Tovar, F. Vasques, "Guaranteeing real-time message deadlines in PROFIBUS networks", **Proceedings of the 10th Euromicro Workshop on Real-Time Systems**, 17-19 June 1998, pp. 79 - 86

[13] H. Kopetz, W. Elmenreich, C. Mack, "A comparison of LIN and TTP/A", **Proceedings of the 2000 IEEE International Workshop on Factory Communication Systems**, 2000, pp. 99 - 107

[14] C. Han, K. Shin, "Real-time communication in FieldBus multiaccess networks", **Proceedings of IEEE RTAS'95**, 1995, pp. 86-95

[15] R. Yavatkar, P. Pai, R. Finkel, "A Reservation-based CSMA Protocol for Integrated Manufacturing Networks", **IEEE Transactions on Systems, Man and Cybernetics** 24:8, August 1994, pp. 1247 - 1258

[16] C.C. Chou, K.G. Shin, "Statistical real-time channels on multi-access bus networks", **IEEE Trans. on Parallel and Distributed Systems**, vol.8, no.8, 1997, pp. 769-780

[17] M.A. Livani, J. Kaiser, "A Symmetric MAC Protocol for CSMA Busses in Dynamic Distributed Real-time Systems", **Proceedings of the Second Int. Conference on Parallel Computing Systems (PCS'99)**, Ensenada, Baja California, Mexico, Aug. 1999.

[18] P. Nain, N.D. Georganas, W.J. Stewart, "Analysis of a hybrid multiple access protocol with free access of new arrivals during conflict resolution", **IEEE Transactions on Communications**, Volume: 36 , Issue: 7 , July 1988 pp. 806 - 815

[19] "Air Interface for Fixed Broadband Wireless Access Systems", **IEEE Standard 802.16**, IEEE Inc., New York, USA, 2002

[20] K. Etschberger, **Controller Area Network (CAN) Basics, Protocols, Chips and Applications**, Ixxat, ISBN: 3-00-007376-0, 2001

[21] H.C.B Chan, V.C.M. Leung, "Credit-based dynamic reservation integrated services multiple access (DRISMA) for wireless ATM networks", **Proceedings of the 2nd Wireless Communications and Networking Conference WCNC**, 23-28 Sept. 2000, pp. 1199 - 1203

[22] P. Brooks, "Ethernet/IP-industrial protocol", **Proceedings of the 8th IEEE International Conference on Emerging Technologies and Factory Automation**, 2001, pp. 505 - 514

[23] **Ethernet Powerlink Whitepaper V5**, BERNECKER + RAINER Industrie-Elektronik Ges.m.b.H., 2002

[24] C. Venkatramani, T. Chiueh: "Design, Implementation, and Evaluation of A Software-Driven Real-Time Ethernet Protocol", **ACM SIGCOMM 95**, 1995