♦ Design Tools for Transparent Optical Networks

Chandra Chekuri, Paul Claisse, René-Jean Essiambre, Steven Fortune, Daniel C. Kilper, Wonsuck Lee, Nachi K. Nithi, Iraj Saniee, Bruce Shepherd, Christopher A. White, Gordon Wilfong, and Lisa Zhang

Optical technology promises to revolutionize data networking by providing enormous bandwidth for data transport at minimal cost. A key to cost reduction is to increase transparency, that is, to keep a data stream encoded as an optical signal for as long as possible. Wavelength switching increases transparency by allowing different data streams, each encoded in a different wavelength of light, to be independently routed through an optical network. We discuss Bell Labs-developed software tools that help design wavelength-switched optical networks. The software tools simultaneously minimize the cost of the designed network, reduce the time and cost to perform the design, and ensure compliance with engineering constraints. The tools span three levels of abstraction, from routing and reconfigurable add/drop multiplexer (ROADM) choice, to span engineering, to power dynamics simulation. Each level represents a different tradeoff between design scope and level of detail. For each class of tool, we briefly describe design philosophy, algorithms, performance, and resulting value for Lucent's customers. © 2006 Lucent Technologies Inc.

Introduction

Optical network technology provides prodigious capacity to transport data [1]. A single optical fiber can carry traffic measured in trillions of bits per second, orders of magnitude more than can be carried by electrical cables or wireless communications. Together with outer layers of packet-switched routers and wireless access, optical networks can provide the data communications infrastructure required by society for the foreseeable future.

The current generation of optical networking, as exemplified by Lucent's LambdaXtreme® Transport System and Metropolis® MetroEON Enhanced Optical Networking products, is based on three technological pillars. First, dense wavelength division multiplexing (DWDM) allows many different independently modulated wavelengths (colors) of light to be carried together on a single optical fiber. Depending upon technology, a fiber can have 32 to 160 different wavelengths, each carrying between 10 and 40 billion bits per second. The second pillar is broadband optical amplification, which allows many wavelengths to be amplified simultaneously without requiring electrical regeneration. The third and most recent pillar is wavelength-granularity optical switching. With such switching, an optical

Bell Labs Technical Journal 11(2), 129–143 (2006) © 2006 Lucent Technologies Inc. • DOI: 10.1002/bltj.20165

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

network can route a single-wavelength lightpath from an arbitrary source network element to an arbitrary destination network element, with no electronic processing required at any intermediate network element.

Optical networks can span hundreds to thousands of kilometers and involve scores to thousands of different network elements. Their design and analysis requires expertise at many levels, from the physics described by the nonlinear Schrödinger equation to the combinatorial solution of NP-hard [7] optimization problems. Optical networks can also require significant investment, ranging from millions to hundreds of millions of dollars. For all these reasons, computer tools are essential to design, optimize, simulate, and operate optical networks. This paper surveys some of the optical design tools developed at Bell Labs over the past few years. These tools address several aspects of optical network design, from helping Lucent Technologies provide optimized bids to customers, to providing a detailed selection of optical components, to planning wavelength growth in operating networks. For lack of space, not all relevant tools are described, nor is any tool covered in much detail. However, the paper does provide a good sample of the computer science expertise required for the design of optical networks.

Panel 1. Abbreviations, Acronyms, and Terms

ATOM—A transparent optical mesh
CROME—Customer's Route via OSNR Model Evolution
DCM—Dispersion compensation module
DGEF—Dynamic gain equalization filter
DWDM—Dense wavelength division multiplexing
MADCAP—MetroEON Amplifier and Dispersion Compensation Assignment and Placement
MNDT—Mesh network design tool
NP—Nondeterministic polynomial time
OA—Optical amplifiers
OSNR—Optical signal-to-noise ratio
OT—Optical transponder
ROADM—Reconfigurable optical add/drop multiplexer
SSMF—Standard single mode fiber

Optical Components

The basic network element in a transparent optical network is a reconfigurable optical add-drop multiplexer (ROADM). Conceptually, a ROADM has one or more arms, an optical switching fabric, and a set of slots for optical transponders (OTs). As shown in **Figure 1**, each arm consists of an optical demultiplexer (trapezoid), with a connection to an external optical fiber on one side and a connection to the switching fabric (square) on the other side. The demultiplexer separates the light from the fiber into its constituent wavelengths. The switching fabric can route each wavelength arbitrarily, either from one fiber to another or from a fiber to an OT slot. Usually the routing is under software control (i.e., reconfigurable), though with some technologies it is effected by a patch panel.

Throughout, we assume that data paths are bidirectional. Hence, for example, if the optical fabric routes a particular wavelength from one fiber to a second, then the same wavelength of light is routed from the second fiber to the first. In a physical implementation of a ROADM, this typically implies that there are distinct data paths in each direction; to simplify the discussion, this duplication is left implicit.

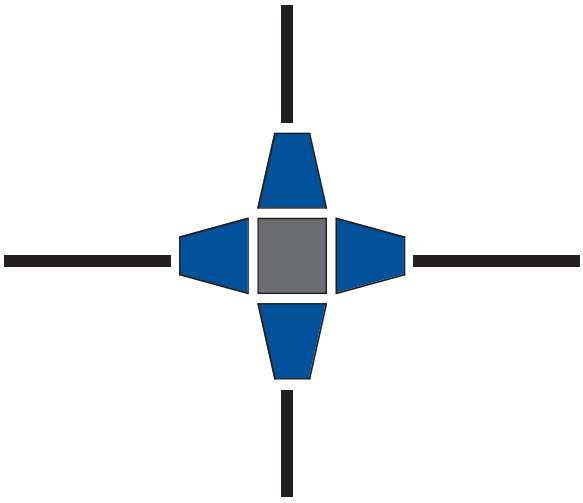

ROADM—Reconfigurable optical add/drop multiplexer

Figure 1.
A 4-arm ROADM.

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

An OT converts an external data signal entering the network into a modulated light source at a particular wavelength and similarly converts a modulated wavelength into a data signal leaving the network. The cost of an OT is usually a significant fraction of the base cost of a ROADM, e.g., 5–10%. Hence if only a few wavelengths are used, the base cost of a ROADM dominates; if most of the wavelengths are used, OT cost dominates.

At the highest level, an optical network is built by connecting ROADMs with optical fibers, populating the ROADMs with OTs, configuring the optical switching fabric of each ROADM, and adding optical amplification to the fiber routes. Then the network provides a set of single-wavelength light paths that start at an OT, flow through various ROADMs, and end at another OT.

Depending upon the technology, the number of arms of a ROADM is limited, from a minimum of one to a maximum of perhaps six or eight. If the number of arms is at most two, then the optical network consists of one or more line systems. Each line system is either a sequence of degree-2 ROADMs arranged into a ring, or a line, starting at an end terminal (a ROADM with one arm), passing through a sequence of degree-2 ROADMs, and ending at an end terminal. If there are several line systems in the network, satisfying a client demand may require several lightpaths, each within a single line system; OTs are needed at the beginning and end of each lightpath.

A mesh network is built with ROADMs with more than two arms. This allows the fiber connectivity to be more complicated than a line or ring, and in general there can be several paths between one ROADM and another. However, even a mesh network may decompose into several components, with no transparent light path possible from one component to another.

Impairments of Optical Data Signals

As a modulated light source traverses the network, it undergoes various impairments that, if not managed, cause the modulating data to be lost. One of the most basic impairments is optical attenuation, which is a decrease in power as a signal traverses a fiber. Optical amplifiers (OAs) come in two common types, Raman and Erbium-doped fiber amplifiers; each amplifies all the wavelengths on a fiber simultaneously.

A second basic impairment is the optical signal-to-noise ratio (OSNR). Any launched light source has a certain noise level. Optical amplification increases both the incoming signal and incoming noise by the same ratio, but also adds additional noise, hence decreasing OSNR.

The third basic impairment, chromatic dispersion (also called group-velocity dispersion), results from slight wavelength dependence of the speed of light within a fiber, causing pulses to spread as they travel. Chromatic dispersion can be compensated with a dispersion compensation module (DCM), which essentially consists of rolls of specially designed fiber with reverse wavelength dependence of light speed. Other impairments, e.g., polarization mode dispersion, or power tilt, are generally less significant.

An optical network must include sufficient optical amplifiers and dispersion compensation modules to ensure that no data is lost to impairments; in particular each light path must exceed the minimum required signal strength and OSNR and not exceed the maximum allowed dispersion. Placing OAs and DCMs is not trivial, because of complex interactions among impairments and components. For example, different optical amplifiers may have different tradeoffs between amplification and OSNR. The OSNR degradation introduced by an optical amplifier usually increases with absolute amplification, suggesting frequent use of low-gain amplifiers; however, placement of optical amplifiers may be constrained by the availability of huts along a fiber path.

Different design choices will result in different requirements with respect to optical impairments, and hence different capabilities. For example, Lucent's MetroEON system is designed for metro and regional networks and provides up to 32 wavelengths at 10 Gbps with an optical reach of 600 km. Lucent's LambdaXtreme system is designed for ultra long reach, 1000 km to 4000 km transmission, and ultra high capacity, either 64 wavelengths at 40 Gbps each or 128 wavelengths at 10 Gbps each.

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

Three Design Problems

Transparent optical network design is too complex to be accomplished by a single holistic computer design tool. Instead we split network design into three phases—routing and ROADM choice, span engineering, and power dynamics simulation—each accomplished by a single class of design tools.

Routing and ROADM choice requires as input a set of client traffic demands and a fiber network, that is, a set of central office locations to place optical network elements and the existing fiber connections between the office locations. The output is a set of configured ROADMs that create the light paths required to satisfy routing demands. The cost of each design includes the ROADMs, OTs, OAs, and DCMs; the latter two are estimated using simplified rules.

Span engineering requires the design resulting from the first phase. A ROADM within a central office typically contains an OA in the output stage, and possibly also a DCM. Along each fiber path there is a set of additional hut locations where OAs and DCMs can be placed. The span engineering phase chooses the OAs and DCMs to correctly mitigate all optical impairments, at minimal cost. This phase is made computationally feasible only by using appropriate simple algebraic models of optical fiber, OAs, DCMs, and impairments, rather than a fully detailed first-principles optical simulation.

Power dynamics simulation requires the detailed configuration produced by the first two phases. It simulates the change in overall amplifier and channel power levels as new wavelengths are added or in response to transient effects such as a fiber cut.

Routing and OADM Choice

We next describe a mesh network design tool (MNDT), and briefly a predecessor tool, Ocube [5]. MNDT performs high-level network design; that is, MNDT chooses the lightpaths to satisfy a given set of a traffic demands in a given fiber network. The fiber network is presented as a graph with nodes representing central office locations and links representing existing fiber paths connecting the central offices. Each traffic demand represents bandwidth equivalent to a single wavelength from one central office to another, possibly with a protection requirement. Although not discussed here, traffic demands can also be given as fractions of a wavelength, requiring grooming, or as "ring demands," which provide the transport required for SONET rings. MNDT has auxiliary parameters that specify the number of arms in a ROADM, the fiber capacity in terms of wavelength, as well as simplified OSNR-based rules for optical reach, and equipment costs.

MNDT output specifies the ROADMs, their configurations, the fibers, and the OTs to provide the light paths required to satisfy each demand. It is possible that a demand may require several consecutive light paths, either because of insufficient transparency, necessity of optical regeneration, or wavelength conversion. In high-demand scenarios, it may be necessary to use several fibers on each link. The objective of the tool is to minimize equipment cost.

This optimization problem is complex, and certainly NP-hard [7]. With a predecessor tool Ocube [5], we first experimented with an integer programming [6] formulation of the problem. This formulation required a large number of integer variables describing routes and ROADM configurations, and was infeasible for all but tiny networks. Furthermore, some constraints such as optical reach were hard to model accurately. With Ocube, we instead developed a partitioning of the overall problem into subproblems with effective heuristics. This approach scaled well to reasonable network sizes; for example Ocube was used for networks from 20 nodes and 30 links to 280 nodes and 340 links, with computation time ranging from a few seconds to a few hours.

MNDT is based on a similar partitioning into three subproblems: routing, ROADM configuration, and light path assignment. For each demand, the routing module chooses a link path, i.e., the sequence of links traversed to satisfy the demand, without specifying a specific fiber or wavelength. The ROADM configuration module specifies how fibers are connected via ROADMs at each node. For each demand, the light path module specializes the link path into a sequence of light paths by specifying the fiber and wavelength for each link along the path.

Routing Module

The routing module determines a link path for each demand with the goal of minimizing estimated

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

network cost, where estimated network cost is the sum of estimated ROADM and OT costs. The number of fibers required for a graph link uv is $f(uv) = \lceil L(uv)/\mu \rceil$, where $L(uv)$ is the number of link paths traversing uv and μ is the capacity of a fiber in wavelengths. ROADM cost can be estimated from the number of fibers on all links, since the cost of a ROADM is approximately proportional to the number of its arms, and each fiber requires two ROADM arms. The OT count required by a link path is estimated by respecting optical reach but ignoring wavelength assignment and lack of transparency. In particular, the link path is split into a minimal set of subpaths so that the total OSNR on each subpath meets the OSNR bound; the number of estimated OTs is then twice the number of subpaths.

Total estimated network cost is reduced using a bypass heuristic that attempts to remove unnecessary fibers. We assume that the initial link route of a demand is either given as input or else follows the shortest path between the source and the destination nodes. The bypass heuristic examines one fiber at a time. When fiber e is bypassed, we reroute all the demands that were using e and eliminate e if the new routing reduces the estimated total cost. The rerouting is based on a weighting function that combines a penalty for filling up a fiber on a link (since further rerouting might require a new fiber on the link) with the estimated OT cost of the path.

ROADM Configuration Module

The ROADM module determines the connectivity of fibers to ROADMs. Every fiber incident to a node must be connected to a ROADM, and to maximize transparency, it is desirable to use ROADMS with as many arms as possible. However, because of technology constraints, the number of arms may be bounded, and hence there is a choice to be made of the actual configuration. For example, if there are five fibers incident to a node, and available ROADMs have up to 4 arms, then at the node there are 15 possible configurations: five configurations of one 4-arm ROADM and one 1-arm ROADM, and ten configurations of one 3-arm and one 2-arm ROADM.

We use two heuristics, MaxReduction and MaxThru. Both heuristics use a local greedy approach that examine nodes one at a time and choose the "best" configuration locally for each node. MaxReduction begins with 1-arm ROADMs only at each node. For each node u, it iterates through all possible configurations at u and for each configuration it calls the light path module to compute the total network cost. MaxReduction then assigns to u the configuration that reduces the cost the most. The configuration at u is now fixed, and the next node is examined. The light path module can be computationally expensive, and use of MaxReduction to calculate "best" configuration can therefore be expensive as well.

MaxThru offers a faster alternative. The through flow at a node u is the number of demands passing through u that can flow transparently through u, i.e., that do not have to switch from one ROADM to another. MaxThru iterates through all possible configurations at u and chooses the one that maximizes the through flow. This computation is straightforward if each link incident at u has only one fiber. If some link has more than one fiber, then the module must first make an arbitrary, temporary assignment of fibers to each path using the link. The final assignment is made in the light path module.

Since MaxThru simply sums up the through flow at each node instead of going through the light path module, it is faster than MaxReduction. However, through flow is only an indication of the number of required OTs, so MaxThru does not minimize the exact objective function whereas MaxReduction does. Therefore, MaxReduction often yields better designs. Hence there is a time-performance tradeoff between MaxThru and MaxReduction.

Light Path Module

Once the ROADM configuration at each node is fixed, the light path module finds the sequence of light paths required to satisfy each demand. For a particular demand, the light path sequence follows the link path determined by the routing module; however if there is more than one fiber on a link, the light path module must choose which fiber to use on that link, and it must also choose the wavelength on the fiber. The light path module also enforces reach limitations on light paths, based on an estimated OSNR contributed by

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

each link. An OT is required at the beginning and end of each light path; thus OTs are required when switching from one ROADM to another, when changing wavelengths, and if the reach limit is exceeded. The overall goal is to minimize the total number of OTs.

The light path module orders the demands by their link count, processing those with the most links first, since demands with long paths are more likely to have a wavelength conflict. For each demand, a dynamic programming algorithm finds the minimum-OT sequence of light paths that follow the link path. The dynamic programming algorithm conceptually splits each fiber on a link into a number of parallel links, one for each unused wavelength. The algorithm traverses the link path, computing for each fiber-wavelength pair at a link the minimum number of OTs required to route the demand to that pair. For example, if fiber f follows fiber e along the link path, and e and f are both connected to the same ROADM, the number of OTs required to reach a wavelength on f is the same as the number required to reach the same wavelength on e. However, additional OTs are required for a wavelength change, a switch from one ROADM to another, or if the OSNR reach limit is attained.

The dynamic programming approach can be time-consuming if the number of fibers, wavelengths and demands is large. A preprocessing step greedily tries to find a wavelength w that is available for the entire link route of demand d. If such a wavelength exists, then we assign w to d. Otherwise, we set the demand aside. We then run the dynamic programming procedure outlined above on all demands set aside. Typically the vast majority of the demands are processed during the greedy preprocessing step, speeding up the entire module.

If the network consists only of linear line systems, from end terminal through 2-arm ROADMs to end-terminal, then fiber assignment is still required, but wavelength assignment is in fact trivial. This follows from a classic interval-graph coloring result: if at most L subintervals of a line segment contain any point of the segment, then the subintervals can be colored with at most L colors so that overlapping subintervals have distinct colors. This coloring can be found by ordering the subintervals by their left points and assigning each subinterval the first color that does not conflict with any overlapping subinterval. This approach was used by Ocube, which handled only linear line systems.

Results

Ocube has been used to respond to dozens of customer requests for network design. In one early case, Ocube saved 17% (or a nominal $35 million) over an earlier manual design for a customer network; computing time was a matter of minutes as compared to the weeks required for the manual design. Subsequent designs were not performed by hand. The designs produced by MNDT for mesh networks have been compared with several manual designs performed by experts; in each case the MNDT design was the same cost or cheaper, and required only a fraction of the total design time.

Span Engineering

Span engineering involves choosing OAs and DCMs for each fiber link to ensure that all optical impairments are sufficiently minimized. A fiber link starts and ends at a central office containing a ROADM and consists of a sequence of fiber spans of known loss and length. Intermediate endpoints of spans are huts. OAs and DCMs can be placed either in huts or in central offices. The overall goal is to minimize equipment cost. For this section we consider the following impairments:

- *Power:* Optical amplifiers typically attempt to adjust amplification to maintain a constant output power, independent of input power. A low-gain optical amplifier might have a maximum of 20 dB of amplification with a target output of -5 to -1 dBm; a high-gain amplifier might amplify up to 33 dB with a target output of 4 to 5 dBm. Power loss in a fiber is typically .2 dB per kilometer, hence optical amplification is required every 100 km or so.
- *OSNR:* After the first OA after an OT, the OSNR might be 55 dB, and as a rough rule, each subsequent OA might reduce OSNR by 3 dB. A typical OSNR required at an OT might be 17 to 19 dB for a 10 Gbps signal at an uncorrected bit-error rate of 10^{-3}. However, the actual OSNR degradation introduced by an OA is a complicated function

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

of the input power, the type of the OA, and the actual gain.

- *Dispersion:* A typical standard single mode fiber (SSMF) introduces 17 ps/nm/km of chromatic dispersion—that is, the time to traverse a km of fiber varies by 17 ps per nm of wavelength. For the purposes here, dispersion can be viewed as a linear function of the length of fiber. A DCM negates the effect of dispersion in the fiber, and thus a DCM is often measured in km of fiber; i.e., a 10 km DCM negates the chromatic dispersion in 10 km of SSMF.
- *Ripple:* An optical amplifier can introduce gain ripple, that is, the amplifier induces a variation in the power level of different wavelengths. Like OSNR, gain ripple is a function of gain and OA type; an OA might introduce a fraction of a dB of gain ripple, with an overall allowed ripple of 3 dB. A special type of OA with a dynamic gain equalization filter (DGEF) reduces gain ripple but also decreases OSNR more than other amplifiers.

CROME and LambdaXtreme

CROME (Customer's Route via OSNR Model Evolution) [2, 3, 4] is a span engineering design tool for Lucent's LambdaXtreme system. CROME assumes that the optical network consists of a set of line systems, where each line system starts at an end terminal, traverses an alternating sequence of links and 2-arm ROADMs, and ends with a link to a final end terminal. Span engineering for each line system can be done separately.

The LambdaXtreme system has some hardware characteristics that simplify span engineering, at least relative to the MADCAP (MetroEON Amplifier and Dispersion Compensation Assignment and Placement) tool discussed below. First, LambdaXtreme ROADMs require that any dispersion within a link must be completely compensated within the link, to within a few hundred ps/nm. This implies that DCM placement can and should be done independently for each link. Second, as a long haul system, LambdaXtreme requires the use of DGEF amplifiers in order to reduce ripple. However, this implies that ripple is in fact relatively small, so individual wavelengths will have approximately the same power, and it suffices to model a nominal wavelength power and ripple.

Both DCM placement and OA placement can be chosen using a dynamic programming algorithm. In fact the two problems decouple, with DCM placement solved first. We sketch the OA algorithm, as it is a bit more involved. For the purposes of OA placement, a DCM just acts as additional span loss; of course, the amount of compensation determines the extra loss. The OA algorithm places OAs to guarantee that any light path that traverses the entire line system meets power, OSNR, and ripple requirements; this guarantees that any light path that traverses only a part of the line system also meets the requirements. The LambdaXtreme Transport System has three kinds of OAs; for the algorithm described next, a fiber splice acts as a fourth possible kind of "optical amplifier."

The dynamic programming algorithm traverses the line system, hut by hut from one end to the other. At each hut, the algorithm computes a table indexed by (power, OSNR, ripple) triples, with discrete values for these quantities chosen to make an appropriate tradeoff between table size and discretization error. Conceptually, the table stores the cheapest sequence of OAs that will attain power, OSNR, and ripple values at least as good as the table index. In fact, the table stores just the optimal cost plus the final OA of the sequence, together with a pointer to a table entry at the previous hut. Assume that the complete table has been computed at the previous hut. The table entry for a current (power, OSNR, ripple) triple is obtained by considering, for each possible OA and each previous table entry, whether the combination of previous (power, OSNR, ripple) triple plus OA plus span loss will attain the current (power, OSNR, ripple). The cheapest choice is recorded as the table entry. After the entire line system is traversed, the cheapest OA sequence can be reconstructed from the cheapest table entry at the final hut that satisfies the requisite power, OSNR, and ripple bounds.

The actual dynamic program used by CROME is somewhat more complex than described above. For example, though the optical amplifiers must be chosen separately in each direction along the line system, the

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

two directions are coupled since splice locations must be chosen identically. Also, various methods can prune the computation required for each table entry.

Flexible Dispersion Maps and MADCAP

MADCAP (MetroEON Amplifier and Dispersion Compensation Assignment and Placement) is a span-engineering design tool for Lucent's MetroEON system. MetroEON is similar to LambdaXtreme: It has 2-arm ROADMs, although rings as well as lines are allowed, and it has a choice of multiple OA types and allowed values of dispersion compensation. The MetroEON system differs from LambdaXtreme in that it uses an aggressive "flexible dispersion map" design methodology to reduce the number of required OAs and to consolidate DCMs. This design methodology allows the different wavelengths traversing a fiber to have different dispersion levels. The design methodology also allows a tradeoff between dispersion and OSNR at an OT; for example, a light path with optimum dispersion value can tolerate a reduced OSNR relative to a path with high dispersion.

This design methodology introduces several complications to span engineering. First, power and OSNR must be modeled on a per-wavelength basis, rather than in aggregate. Second, it is not possible to decouple DCM and OA selection, since the choice of one may restrict the possible choices of the other. In consequence, the dynamic programming approach of CROME does not apply, since the state space (table size) would be much too large.

Instead, the actual MADCAP tool is designed with two layers, an inner simulator and an outer search algorithm. The inner simulator can test whether a specific configuration, i.e., a placement of OAs and DCMs at huts, satisfies all the engineering rules. This simulator is a bit complex: it simulates dispersion, power, and OSNR evolution of multiple wavelengths independently, taking into account that the wavelengths can enter and exit the line system at different points. It also models the Raman effect, where the power from one wavelength is transferred to another, and includes an algebraic model of the tradeoff between OSNR and dispersion.

Two different search algorithms try to find a minimum-cost configuration. One search algorithm is a randomized iterative algorithm based on Simulated Evolution [9]. This algorithm starts with a plausible initial configuration. A per-hut fitness function is computed based on equipment cost and whether local engineering constraints are satisfied. At each iteration, the equipment in each hut is changed in a random way, with unfitness increasing the probability of change. The iterations stop at a set count, or when all engineering constraints are satisfied and the equipment cost no longer decreases substantially.

A second search algorithm tries to search the entire configuration space of OA and DCM placements. This configuration space is enormous—there can be tens or hundreds of hut locations, each with one of three OA types and roughly twelve different DCM types. Two algorithmic ideas prune the search space. One is to restrict the search to just the space of OA configurations, and from an OA configuration deterministically obtain the requisite DCMs. This latter step is nontrivial. Adding DCMs increases loss, hence requires more amplification, and hence increases OSNR. There are actually two subalgorithms for DCM placement: a relatively efficient one to maximize OSNR, and a slower one to minimize cost; the former is used as an estimate and the latter in a final optimization phase. The second algorithmic idea is a form of dynamic programming: find all the satisfactory OA configurations that satisfy the engineering constraints for each link individually, then combine these configurations for pairs of adjacent links, then for triples of links, and other operations.

Both algorithms actually produce similar results; the deterministic algorithm is generally faster than the random iterative algorithm; however it has more dependence on the size of the configuration space. In the end, there were two versions of the MetroEON hardware, one with variable-gain OAs and one without. Because of the larger search space with variable gain OAs, the random iterative algorithm was used for the former version and the deterministic algorithm used for the latter version.

Results

Figure 2 plots the relative cost of a MADCAP design versus a manual design using simple engineering rules. MADCAP costs appear in darker color. The

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

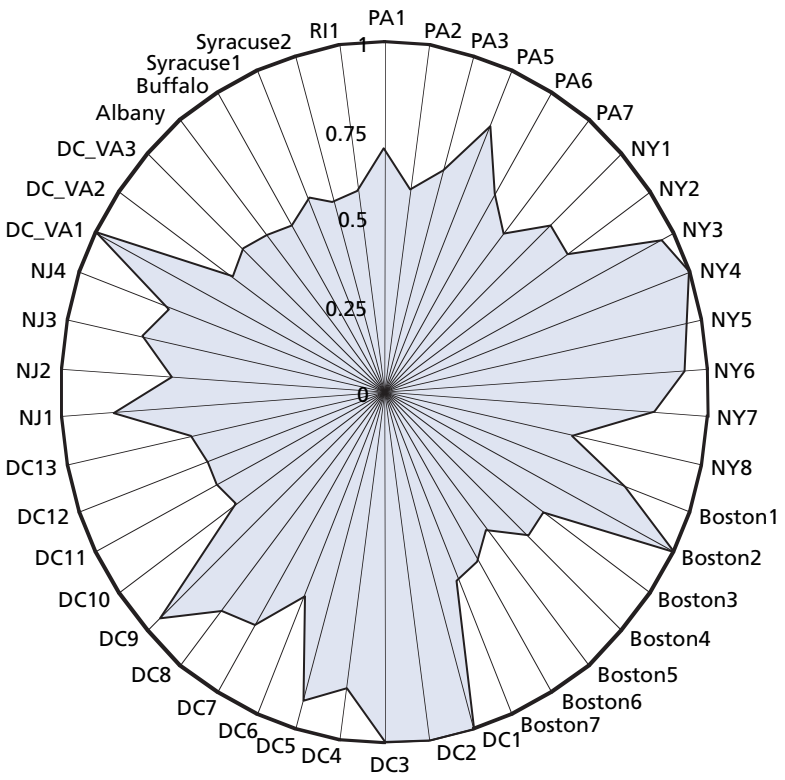

Figure 2.
The relative costs of various networks designed with MADCAP versus manual design.

names, e.g., PA1, denote various actual and synthetic networks. Every network is a single ring. While the benefits of the MADCAP design vary widely from one network to another, in many cases the MADCAP design is substantially the cheapest.

Power Dynamics Simulation

The size and complexity of large-scale optical networks often requires use of hardware testbeds or small prototype networks to infer the physical behavior of larger networks. Simulation provides a complementary method that augments such experimentation with the ability to scale to networks that are too large or too costly to consider directly.

One question that requires extensive simulation concerns the method of achieving and maintaining steady state optical power distribution within a mesh network; in particular, understanding how a change in power at one node (from adding/deleting channels or adjusting amplifier properties) causes power changes at other nodes of the network. In a transparent network, these power changes can traverse the entire network potentially causing significant impairments to transmission. Understanding these power changes is crucial to designing a stable optical mesh network.

An ideal network simulator for addressing this question would solve the time-dependent differential equations for signal propagation and amplification as well as include an implementation of all control functions of the system [8]. Unfortunately, in the context of large mesh networks, this type of all-encompassing simulation quickly becomes computationally intractable. The ATOM (A Transparent Optical Mesh) simulator produces a computationally feasible simulation by addressing only the details relevant for understanding the basic behavior of a transparent optical mesh. ATOM contains a flexible simulation architecture that enables one to shift the focus of the simulation from complex

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

mesh topologies with simple physics to simple topologies with detailed physics. Within a single simulation, portions of the network surrounding a node of interest can implement the detailed physics while nodes in the periphery utilize simple physical models.

For many aspects of the mesh, the topology and traffic completely dominate the behavior of the system. A typical nationwide network can have hundreds of ROADMs leading to the simulation of thousands of interconnected nodes. Typical physical layer simulators [8, 10, 11], which only treat optical line systems, cannot be used to understand the potential feedback effects arising from the mesh topology. Further, the large number of nodes often leads to intractable computational problems. For these reasons, it was important for the ATOM simulator to allow arbitrary connectivity between nodes in a computationally tractable manner.

To understand transparent optical mesh networks through simulation, one needs to apply flexibility to address and model the network topology, traffic, and amplifier physics. In contrast with conventional line systems, which typically have a single unique route between two nodes, a mesh topology has multiple paths connecting any two nodes. The existence of multiple paths enables the construction of cycles giving rise to the possibility of feedback effects within the network. Obviously, nodes that drop all incoming demands or links that transport no traffic will effectively break this connectivity. For this reason, sets of nodes that route entirely disjoint sets of demands will not interact and thus will behave as separate mesh networks. Conversely, the longer the paths of two demands overlap, the greater the amount of coupling introduced between nodes of the system. Amplifier physics provides the final mechanism of coupling. As an artifact of broadband multi-channel amplification, a signal-signal power interaction arises either as a direct physical effect, or as a result of amplifier adjustments required to achieve a desired amplifier gain profile. The real advance of the ATOM simulator concerns the creation of a simulation architecture that exposes the interplay between topology, traffic, and amplifier physics in a computationally tractable manner.

Within this section, we describe the properties of the ATOM simulator for computing the power level physics of the transport layer of transparent optical mesh networks. The simulator allows a detailed treatment of network topology, traffic, and amplifier physics. The flexible, object-oriented design enables one to adjust the focus of the simulator from complex topology with simple physics to detailed physics with a simple topology. We apply the simulator to a simple mesh network and compute the time dependent and steady state reaction of the network to changing the input power levels.

Topology and Traffic

Figure 3 provides an overall view of ATOM simulator components. The ATOM simulator achieves arbitrary connectivity through the concept of a node, shown in Figure 3a. A node object allows an arbitrary number of input and output ports and can represent either a ROADM or a simple repeater. A node contains a routing table that holds wavelength level routing information between input and output ports. Each input port contains a pointer to (representing connection to) an output port of another node of the network. The output ports contain an amplifier (possibly composite) implementing the physics of transport and amplification. In this manner, all routing and physics details remain locally defined within each node. Because operations on a node do not require information from nodes other its nearest neighbors, this locality provides scalability in network computation, control, and storage. This contrasts with methods which instead hold tables of global information (e.g., connectivity or demands) that typically grow nonlinearly with the size of the network.

The computation of steady state power for a node of the network occurs by interrogating a given output port of a node for its power spectrum at a specified time. The routing table defines a set of input ports and channel numbers necessary to compute this output spectrum. We recursively obtain the power spectrum at each input port by querying the connected output ports from other nodes at an earlier time (representing the propagation time between nodes). The process of determining powers at earlier times continues until we arrive at a known power level—either by reaching the steady state power or a previously cached power

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

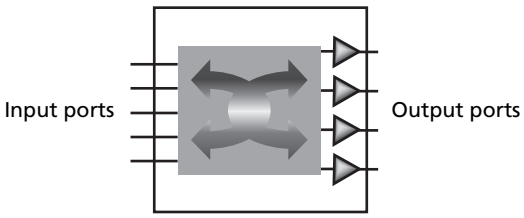

(a) Representation of a simple node.

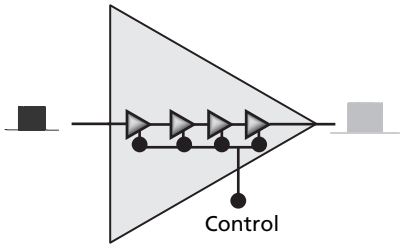

(b) Representation of a composite amplifier.

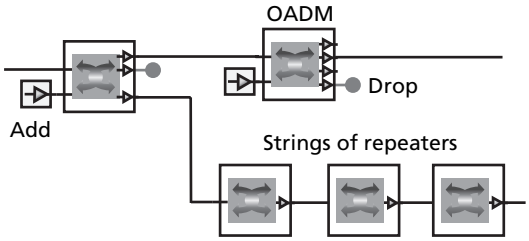

(c) Construction of a network by linking nodes.

OADM—Optical add/drop multiplexer

Figure 3.
Atom simulator components.

computation. In this manner, the power request traverses the network interrogating only those nodes, ports, and times needed to compute a target output port power. This contrasts with conventional simulators that typically would compute the power levels at every point in the network for every time up to the time of interest.

Amplifier Physics

We have designed a hierarchy of amplifier models to reduce the computational load of performing a simulation. By first simulating very simple models within a complex topology, we gain insight into the network using reasonable computational resources. Building on this understanding, we can increase the level of sophistication in the amplifier models either within local portions of the network or within the entire network for a single "hero" computational experiment. Central to the design of the ATOM simulator is the ability to adjust the simulation focus from the topology and traffic to the detailed amplifier physics, enabling rapid progress in understanding the complexities of mesh networks.

Each output port contains an amplifier object, as shown in Figure 3c. This amplifier object takes an input spectrum and produces an output spectrum that represents the amplifier physics and any internal control state of the amplifier. We have implemented typical algebraic amplifier models (linear, constant gain, and saturated) as well as solvers for the detailed physics of Raman and Erbium-based amplifiers. Since amplifier objects can contain other amplifier objects, one can easily construct hybrid multi-stage amplifiers typical of most modern networks, Figure 3b. The object-based implementation means one can easily replace the amplifier object within a node without influencing the topology or traffic. This enables direct discovery of effects related only to the physics of the underlying amplifiers. In addition to the understanding of mesh networks, this object-oriented method provides an ideal platform for the development and testing of new approximate amplifier models.

Power Stability in Optical Mesh Networks

One immediate application of the ATOM simulator concerns the determination of the steady state

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

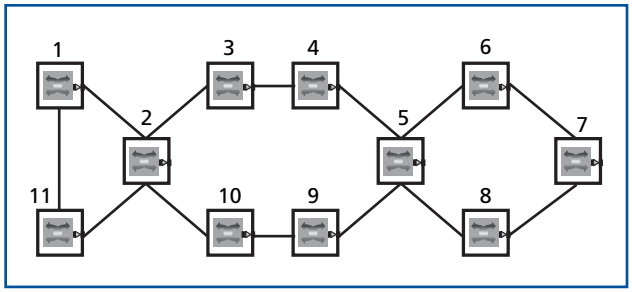

Figure 4.
A simple optical mesh network.

power stability of an optical mesh network. **Figure 4** shows a simple optical mesh network. The question of power stability concerns how changing the power at one node of the system will affect the outputs from the other nodes in the network. This change in power could arise from provisioning a new demand, or from the amplifier adjustments necessary to achieve optimal signal transmission. Practically, these simulations indicate which nodes of the system are most at risk for stability problems. Fundamentally, these simulations show how adjustments in the basic properties of the network topology and traffic lead to the construction of the most stable network possible.

Figure 5 shows a representation of the coupling between input and output nodes for a simple mesh network using ideal saturated amplifiers. These

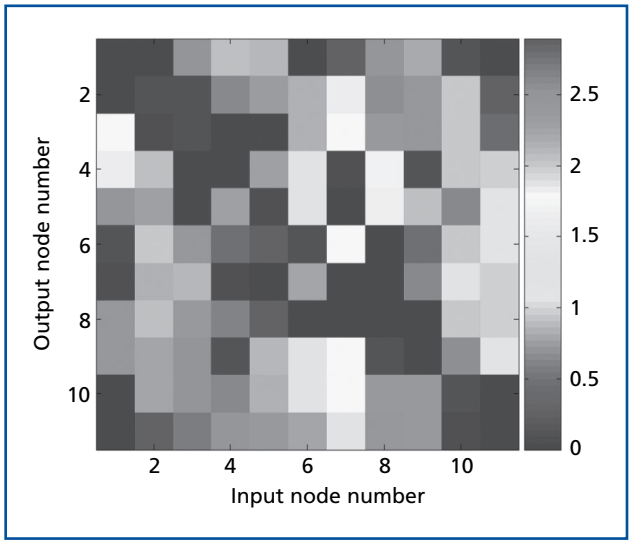

Figure 5.
Coupling between input and output nodes in a simple mesh network.

amplifiers maintain a constant total output power independent of the input channel powers. The amplifier equally scales individual output powers for each channel to maintain a constant total output power. Each column of the grid shows the maximum change in output power for every node of the system (y-axis) for a 3 dB decrease in channel power for all of the channels added at a specified input node (x-axis). Apparent from the plot is the weak interaction between neighboring nodes. This effect arises because the average demand length for this system happens to be four hops. Thus, little or no traffic couples neighboring nodes. Channels dropped at a neighbor will only share a single hop with the newly modified channel powers, thus the impact is much smaller than channels that share multiple hops.

Equally important to the steady state properties described in Figure 5 are the time dependent results for this type of network. We wish to find any uncontrolled feedback effects within the network as well as understand the damping of these feedback effects and the approach of the network to a global steady state power distribution. In this example, we focus primarily on the time scale of network propagation (ms) rather than the time scale of the intrinsic amplifier physics (μs). For this reason, we compute power changes in the local amplifier steady states and neglect any intrinsic time dependence of the amplifiers transitioning between local steady states. Aside from being an excellent approximation due to the relatively large difference in time scales (ms vs. μs) this dramatically decreases the computational effort to obtain time dependent results.

Figure 6 shows the time dependent power changes of the channels dropped at node 5 resulting from changing the power of the channels added at node 5. Because the add and drop channels sets are distinct, the power changes to the drop channels occur only through higher order power transfers. (i.e., interactions between three or more channel groups) The first change in power occurs four hops after the initial event; this corresponds exactly to the smallest cycle of the network that includes node 5. The other power changes occur at a number of hops corresponding to combining various cycle times in the network.

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

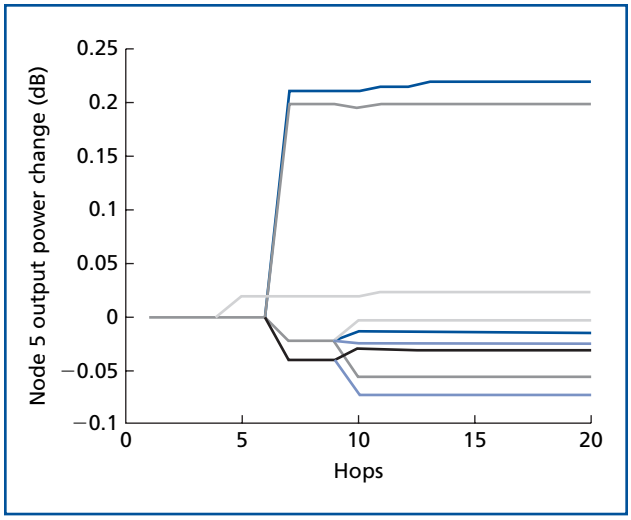

Figure 6.
Network time dependence of the channels dropped as a result of the decrease in input power for channels.

Also of importance, one sees the total power change achieved for a given channel can be larger while the total network adjusts than the final steady state power change for that channel.

Conclusions

Optical networks are sufficiently complicated that manual design is time consuming, error-prone, and suboptimal in cost. Tools such as the ones described here should be viewed as integral parts of the overall optical systems, since these tools can help minimize both the cost of a design as well as the time required to develop it. This methodological improvement allows more design alternatives to be considered while improving the overall quality of the resulting optical networks.

There are many directions for future research. For example, MNDT currently performs single-year designs, while most customer scenarios include growth over several years. Adding incremental design to MNDT substantially complicates the required bookkeeping and also requires examination of tradeoffs between optimizing for first-year or total cost. Similarly, it is a considerable challenge to extend combinatorial span engineering algorithms from line systems to full mesh networks. Finally, a persistent theme for simulations like power dynamics is to populate further the trade-off between accuracy and computational efficiency, for example, to obtain amplifier models that are almost as efficient as algebraic models and almost as accurate as models that are based on detailed physics.

Acknowledgements

Many people have contributed to the work described here. An inexhaustive list includes Jim Ballintine, Jim Benson, Cristina Cannon, Lawrence Cowsar, Dave Einstein, Mohammed El-Sayed, Dan Fishman, Eran Gabber, Bob Feldman, Pradeep Limaye, Mohcene Mezhoudi, Stojan Radic, and Wim Sweldens. We would like to thank the referees for their helpful comments.

References

[1] G. P. Agrawal, Fiber-Optic Communication Systems, John Wiley & Sons, New York, NY, 1997.

[2] Y. Aoki, "Properties of Fiber Raman Amplifiers and Their Applicability to Digital Optical Communication Systems," J. Lightwave Tech., 6:7 (1988), 1225–1239.

[3] C. Chekuri, W. Lee, and L. Zhang, "Optical Line System Configuration Via Dynamic Programming," Proc. Internat. Network Optimization Conf. (INOC '03), (Paris, Fr., 2003).

[4] C. Chekuri, W. Lee, and L. Zhang, "DCM Selection on an Optical Line System," Proc. 11th Internat. Telecommunications Network Strategy and Planning Symposium, (Networks '04), (Vienna, Austria, 2004), 385–390.

[5] S. Fortune, W. Sweldens, and L. Zhang, "Line System Design for DWDM Networks," Proc. 11th Internat. Telecommunications Network Strategy and Planning Symposium (Networks '04), (Vienna, Austria, 2004), 315–319.

[6] R. Fourer, D. M. Gay, and B. W. Kernighan, AMPL: A Modeling Language For Mathematical Programming, Boyd & Fraser Pub. Co., Danvers, MA, 1993.

[7] M. R. Garey and D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman, New York,NY, 1979.

[8] T. K. Ho, T. Salamon, R. W. Freund, C. A. White, B. K. Hillyer, L. C. Cowsar, C. J. Nuzman, and D. C. Kilper, "Simulation of Power Evolution and Control Dynamics in Optical Transport Systems," Bell Labs Tech. J., 10:1 (2005), 119–137.

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

[9] R. M. Kling and P. Banerjee, "ESP: A New Standard Cell Placement Package Using Simulated Evolution," Proc. 24th ACM/IEEE Design Automation Conference, (Miami Beach, Florida, 1987), 60–66.
[10] RSoft Design Group, OptSim Software, <http://www.rsoftdesign.com>.
[11] VPIsystems, VPItransmissionMaker Software, <http://www.vpiphotonics.com>.

(Manuscript approved February 2006)

CHANDRA CHEKURI is a member of technical staff in the Algorithms Research Group within the Mathematical and Algorithmic Sciences Center at Bell Labs in Murray Hill, New Jersey. He completed his undergraduate studies at the Indian Institute of Technology, Madras (now Chennai), and received his Ph.D. in computer science from Stanford University in California. His research interests are in design and analysis of algorithms, combinatorial optimization, theoretical computer science, and their applications. Dr. Chekuri has worked on various tools for Lucent including Ocelot® and CROME, for which he received awards.

PAUL CLAISSE is a member of technical staff in Lucent Technologies' DWDM Optical Network Analysis Group in Holmdel, New Jersey. His current responsibilities include network design and simulation. He has a B.Sc. in physics and electronics from the University of Reading, England, and MBA and Ph.D. degrees from the University of Bradford, England. Dr. Claisse has previously worked for AT&T and Motorola in the areas of semiconductor device design and reliability. He has published 30 papers and has 15 patents issued.

RENÉ-JEAN ESSIAMBRE is a distinguished member of technical staff in Bell Labs' Optical Transport Networks center in Murray Hill, New Jersey. His research focus is on advanced transmission schemes for next and future generations of optical transmission networks. He has a B.Sc. in physics, and M.Sc. in astrophysics, and a Ph.D. in optical physics, all from Laval University in Québec, Canada. Dr. Essiambre received the Engineering Excellence Award from the Optical Society of America in 2005.

STEVEN FORTUNE is technical manager of the Algorithms Research group within the Mathematical and Algorithmic Sciences center at Bell Labs in Murray Hill, New Jersey. He received a Ph.D. in computer science from Cornell University in Ithaca, New York. Dr. Fortune leads a group that performs a broad spectrum of algorithmic research, from fundamental graph-theoretic algorithms to optical and wireless network design tools.

DANIEL C. KILPER is a member of technical staff in the Advanced Photonics Research Department at Bell Labs in Holmdel, New Jersey. He received the B.S. degree in electrical engineering and the B.S. degree in physics from Virginia Polytechnic Institute and State University in Blacksburg, VA, and the M.S. and Ph.D. degrees in physics from The University of Michigan, Ann Arbor. Following his Ph.D. work, he served as a post-doctoral research scientist at the Montana State University Optical Technology Center in the Department of Physics. He became an assistant professor of physics at the University of North Carolina, Charlotte. Dr. Kilper's current research interests include architectures and control systems for transparent and re-configurable optical networks, optical performance monitoring, and the dynamical properties of optical amplifiers.

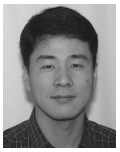

WONSUCK LEE is a member of technical staff in the Complex Systems Analysis and Optimization Department within the Mathematical and Algorithmic Sciences Center at Bell Labs in Murray Hill, New Jersey. His current interests include computational applied mathematics including computational fluid dynamics, computational photonics, numerical methods for differential equations, scientific computing, and parallel computing. He received a Ph.D. in applied mathematics from the State University of New York at Stony Brook. Prior to joining Bell Labs, Dr. Lee worked at IBM Research, and at the Texas Institute of Computational and Applied Mathematics (TICAM). He also spent two years as a postdoctoral fellow at The University of Texas at Austin.

NACHI K. NITHI is a member of technical staff in the Complex Systems Analysis and Optimization department within Bell Labs' Mathematical and Algorithmic Sciences center in Murray Hill, New Jersey. He earned a B.E. in electrical engineering from Madras University, Chennai, India, an M.E. in computer science from Anna

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.

University, Chennai, India, and a Ph.D. in computer science from Colorado State University, Fort Collins, Colorado. Dr. Nithi's interests include tools for optical network switching center design and optimization, software engineering, web technologies, simulations and wireless network monitoring applications.

IRAJ SANIEE is the director of the Complex Systems

Analysis and Optimization Department within the Mathematical and Algorithmic Sciences Center at Bell Labs in Murray Hill, New Jersey. He received his B.A. (Hon.) and M.A. (Hon.) degrees in mathematics and his Ph.D. in operations research and control theory, all from the University of Cambridge. At Bell Labs, the emphasis of research in Dr. Saniee's department has been on network design and mathematical modeling, analysis, and optimization of emerging processes in wireless, optical, and data networks. Recent joint projects with Lucent product groups include new schemes for scheduling and network routing algorithms. His personal research interests are currently in new modes of optical networking, randomized resource allocation, and distributed optimization. Prior to joining Bell Labs, he was a director in the Applied Research Area of Bell Communications Research in Morristown, New Jersey. Dr. Saniee is past chair of the Telecom Section at INFORMS, a member of IEEE, a member of the IFIP Working Group 7.3 on Computer Performance Modeling and Analysis and is currently an associate editor of Operations Research. He led the 1994 runner-up team in the INFORMS' Edelman International Prize in Operations Research and has published numerous articles in IEEE, INFORMS, and SIAM journals and proceedings.

BRUCE SHEPHERD is a member of technical staff in the

Fundamental and Industrial Mathematics group within the Mathematical and Algorithmic Sciences center at Bell Labs in Murray Hill, New Jersey. He earned a B.Sc. in combined math and computer science at the University of Victoria in British Columbia. He went on to obtain a Masters degree in math and a Ph.D. in combinatorics and optimization at the University of Waterloo in Ontario. Dr. Shepherd held a NATO Fellowship in 1990 at CWI, The Netherlands and a Von Humboldt Fellowship at the University of Bonn in 1991. He joined faculty at the London School of Economics (LSE) and held the Nuffield Fellowship for newly appointed lecturers. He was a tenured Reader at LSE when he first joined Bell Labs in 1996.

CHRISTOPHER A. WHITE is a distinguished member of

technical staff in the Complex Systems Analysis and Optimization Department within the Mathematical and Algorithmic Sciences Center at Bell Labs in Murray Hill, New Jersey. He holds B.S. and M.S. degrees in chemistry from Carnegie Mellon University in Pittsburgh, Pennsylvania and a Ph.D. degree in theoretical quantum chemistry from the University of California in Berkeley, California. His research interests include the development of computational methods for the simulation of interesting physical systems. Example projects have included acoustic scattering in colloidal suspensions, the design of nanostructured materials, and the design and optimization of optical devices. Currently, Dr. White's work focuses on the simulation, control and stability of highly transparent optical mesh networks.

GORDON WILFONG is a distinguished member of

technical staff in the Algorithms Research Group within the Mathematical and Algorithmic Sciences Center at Bell Labs in Murray Hill, New Jersey. He received his B.Sc., First Class Honours, in mathematics from Carleton University in Ottawa, Ontario and his M.S. and Ph.D. in computer science from Cornell University in Ithaca, NY. His major research interests are in algorithm design and analysis. Dr. Wilfong's current area of focus is networking including protocol analysis, routing, scheduling and analyzing game theoretic network models. He holds patents in a wide range of areas including protocol design, optical networking, handwriting recognition, security, document analysis and cross-connect design.

LISA ZHANG is a member of technical staff in the

Algorithms Research Group within the Mathematical and Algorithmic Sciences Center at Bell Labs in Murray Hill, New Jersey. Her research broadly concerns algorithmic and complexity issues of networking, with a focus on design and optimization, routing and scheduling protocols, and stability and Quality-of-Service in wired and wireless networks. Dr. Zhang received a B.A. in mathematics from Wellesley College in Wellesley, Massachusetts and a Ph.D. in theory of computation from the Massachusetts Institute of Technology in Cambridge. ◆

Authorized licensed use limited to: University of Illinois. Downloaded on September 30,2025 at 18:21:10 UTC from IEEE Xplore. Restrictions apply.