Managing the Future Internet through Intelligent In-Network Substrates

Marinos Charalambides, and George Pavlou, University College London
Paris Flegkas, University of Thessaly
Ning Wang, University of Surrey
Daphne Tuncer, University College London

Abstract

The current Internet has been founded on the architectural premise of a simple network service used to interconnect relatively intelligent end systems. While this simplicity allowed it to reach an impressive scale, the predictive manner in which ISP networks are currently planned and configured through external management systems and the uniform treatment of all traffic are hampering its use as a unifying multi-service network. The future Internet will need to be more intelligent and adaptive, optimizing continuously the use of its resources and recovering from transient problems, faults and attacks without any impact on the demanding services and applications running over it. This article describes an architecture that allows intelligence to be introduced within the network to support sophisticated self-management functionality in a coordinated and controllable manner. The presented approach, based on intelligent substrates, can potentially make the Internet more adaptable, agile, sustainable, and dependable given the requirements of emerging services with highly demanding traffic and rapidly changing locations. We discuss how the proposed framework can be applied to three representative emerging scenarios: dynamic traffic engineering (load balancing across multiple paths); energy efficiency in ISP network infrastructures; and cache management in content-centric networks.

ISP networks today are normally planned and configured in a *predictive* manner through long timescale engineering where the expected traffic demand is calculated from previous usage and a specific network configuration is produced, aiming to optimize resource utilization over the next provisioning period, typically in the order of weeks or even months. Advanced management paradigms with adaptive feedback control-loop functions are still missing, given the current nature of management systems that are *external* to the network and the resulting latency in learning about arising conditions and effecting changes. As such, existing off-line configuration approaches can be well suboptimal in the face of changing or unforeseen user demands and network conditions.

The nature of emerging interactive applications, the planned migration of telecommunication services over the Internet, and also its potential use by safety and mission critical systems, demand better quality, dependability, and resilience. New operational requirements are also being envisaged for the future Internet, for instance energy efficiency where network devices might not always run in their full capacity for power saving purposes. These emerging requirements have introduced brand new challenges for the design of the next-generation network management systems, given that the current management functions that sit *outside* the operating network are, in this context, even more rigid and inefficient in dealing with complex arising conditions. As a result, there is a need for introducing self-management intelligence *within* the network in order to make the latter more flexible and adaptive to changing conditions through feedback closed-loop control solutions.

Embedding management intelligence into the network is a discipline known as *autonomic networking*. This was inspired by IBM's pioneering autonomic computing vision [1], which envisages "systems that can manage themselves given high-level objectives by administrators." While a number of approaches have been proposed related to self-management, most of them are high-level architectures and design of infrastructures [2–5] focusing only partially on the key research issues. This article proposes the concept of intelligent substrates as a key component of self-management architectures. These substrates realize a distributed infrastructure for embedding dedicated management logic in the network in a controllable manner. Separate coexisting substrates are responsible for different management tasks, which collectively achieve sophisticated functionality and are able to adapt and react gracefully to changes through closed-loop control interactions. They can be viewed as a "virtualization" of management functions on top of the physical network infrastructure for different purposes. Each of the in-network substrates relies on common generic self-management functions, but specific substrates with different management tasks/objectives could be flexibly introduced for handling specific emerging management requirements.

The proposed intelligent substrate paradigm extends the concept of autonomic managers in [1] and [2] through the virtualization of management functions and can thus offer high flexibility in realizing *customized* functions, e.g. coordination of decision making processes, for different management tasks within individual substrates. Furthermore, since the management functions of the intelligent substrates are effectively "planted" within network elements, each virtual substrate can be flexibly organized and mapped onto specific physical devices where necessary/ appropriate, according to distinct management objectives. Another key issue in autonomic frameworks is the harmonization of management decisions given the diversity of management operations. Although some communication between autonomic managers/entities has been considered by [3, 4] and [5], the purpose was only to support the same management tasks. The proposed substrates, on the other hand, can also gracefully interact with each other to achieve interrelated management tasks, for instance sharing common network information or performing coordinated decision making. From this point of view, cross-substrate optimization becomes possible for achieving global optimality and configuration stability while covering different management facades.

The purpose of this article is not to present an applied solution for the self-management of a specific problem domain, but rather to propose an architectural framework and principles a self-management architecture should adhere to. As such, we list possible design choices that can be followed when instantiating such a framework by applying it to three representative future Internet scenarios. In the rest of this article we first detail the concept of intelligent substrates and propose an associated self-management overall architecture. We then describe important research challenges associated with the proposed intelligent substrate approach, such as organizational issues and coordinated decision making. We finally validate the design of the proposed framework by applying it to three emerging representative scenarios, namely dynamic traffic engineering, energy efficiency, and in-network cache management.

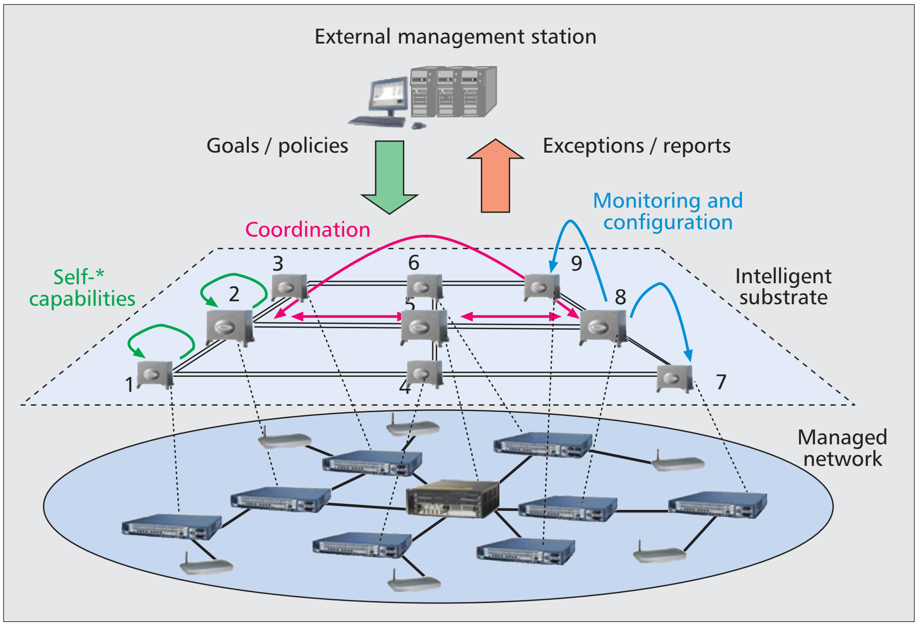

Figure 1. *Intelligent self-management substrate.*

Intelligent Self-Management Substrates and Architecture

The role of the proposed intelligent substrate approach is to form the natural self-managed network environment through parallel and continuous resource management functions, with each substrate supporting a specific network management task by optimizing a specific resource. The term *in-network substrate* emphasizes the fact that while substrates are essential for optimized network operations, they are hidden within the network and, as such, invisible to network users and applications. The term substrate relates to biology which defines it as "the natural environment in which an organism lives."

The intelligent substrate concept is shown in Fig. 1, with only a single substrate shown for illustration purposes. This depicts the physical network at the bottom with the various devices constituting it taking distinct roles in the intelligent substrate and cooperating in order to achieve specific self-management functionality pertinent to this substrate. Although the substrate is depicted above the physical network, it effectively operates *within* the network devices and interacts with other substrates, each dedicated to specific management tasks, for achieving sophisticated self-management functionality with minimum external intervention. The external management system configures initially the operation of the intelligent substrates by setting operational goals or guidelines through high-level policies, setting particular device roles and producing initial resource allocations, e.g. through off-line engineering. It may subsequently adjust those goals as requirements evolve and as it receives exceptions if specific self-management functionality cannot be fulfilled, which may result in goal changing or in a different initial (re-)configuration. Reports are also produced during network operation for recording important self-management decisions, for performance verification, and future planning.

In the specific intelligent substrate example of Fig. 1, the devices numbered 2, 5, and 8 have assumed the role of *regional managers*, with 2 managing 1 & 3, 5 managing 4 & 6, and 8 managing 7 & 9, respectively. These roles may have been preassigned by the external management system, they may be allowed to change dynamically according to substrate operation and self-optimization at runtime, e.g. 1 may substitute 2 for managing 2 & 3, or, in the most fluid case that typically applies to infrastructure-less networks, the local managers are decided by the devices themselves through an adaptive management clustering approach. Manager nodes 2, 5, and 8 cooperate to agree on how they will configure and fine-tune the feedback control loops in their region, including themselves. Node 8 is shown in the figure to monitor and configure its region, i.e. nodes 7 and 9, given the agreed targets with co-managers 2 and 5. Self-management capabilities are shown explicitly for nodes 1 and 2.

Figure 2 proposes the layered internal architecture of a self-managed node, which places management intelligence within the network. The bottom layer includes the hard-engineered control functions, e.g. automatic re-convergence functions for resilience, etc., which are maintained following an evolutionary approach. The same is the case for the management agent layer which still provides a "traditional" interface to the external management system. In addition, an Application Programming Interface (API) provides programmatic access to all the basic resources and the hard engineered con-

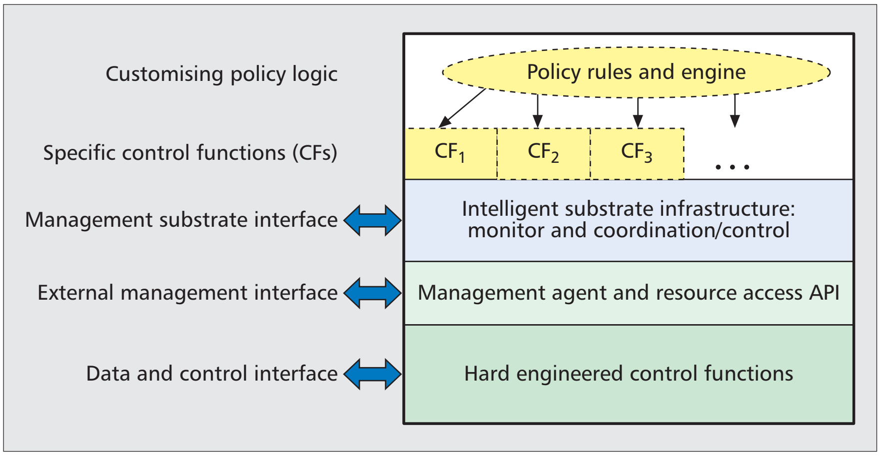

Figure 2. *Self-managed node architecture.*

trol functions of the device. The layer of the intelligent substrate infrastructure includes generic functionality such as monitoring and coordination/control support that exists in every device and is used by all possible substrates. Substrate organization algorithms will be part of this functionality and every device could potentially become a manager in specific intelligent substrates. This substrate infrastructure layer also provides an interface for a node to communicate with other nodes, for example using pub/sub, gossiping or other advanced mechanisms, and also to be accessed by the external management system for control function configuration and for policy introduction. While this layer provides support for coordination/control within specific substrates, it may also provide functionality for coordination among multiple related substrates with distinct management objectives. Specific control functions for each instantiated substrate are depicted as CFi in the figure. Finally, policy decision and enforcement functionality can be used to flexibly configure and extend the relevant control functions through policy logic.

The proposed intelligent substrate approach is evolutionary since it can gracefully coexist with current external management system architectures, although the complexity and functions of the latter are expected to be substantially reduced. On the other hand, it is also revolutionary in the sense that it is introducing radical new capabilities for continuous self-optimization, adaptability and robustness that today's management and control technologies simply cannot deliver.

Intelligent Substrate Organization and Communication

In the proposed in-network management framework described above, the intelligent substrates are formed by dynamically chosen nodes that are responsible for making and enforcing management decisions within individual substrates. Each of these nodes, depending on their assigned role, generates awareness by gathering monitoring information in its area of responsibility and takes management decisions, typically in a collaborative manner with other manager nodes. The key research question is how these nodes are chosen, what criteria drive their selection, and how they are organized into a management structure, i.e. the intelligent substrate. To address this issue, different organizational models may be applied such as hierarchical, distributed, or hybrid approaches, depending on the nature, scale, and characteristics of the managed network as well as the management functionality realized within the substrate. In our proposed framework, substrates of generic management functionality may also exist such as a generic "awareness substrate" responsible for collecting and acquiring information and knowledge that needs to be shared or disseminated to other substrates dedicated to specific management tasks.

Approaches related to substrate organization have been mostly studied in the context of wireless ad hoc networks, investigating node clustering. They focused on "cluster head" selection, the latter being effectively regional managers, by using only topological criteria. In our context, these node selection algorithms should depend on specific criteria and metrics related to the roles assigned to the elected nodes. Criteria such as minimizing monitoring load, achieving consistency in configuration decisions, minimizing the traffic and latency in manager cooperation etc., need to be taken into account. Typical management roles assigned to nodes include those that hold/cache monitoring, aggregated, and processed information, manager nodes with specific control functionality, policy repositories, policy decision points, etc. This node selection process proposed in our framework is a continuous operation that dynamically adapts to changing network and environmental (i.e. context) conditions. When a re-organization of the substrate is performed, relevant substrate functionality is activated in the newly brought in manager nodes. We envisage the situation where nodes cooperatively decide on the most suitable model depending on network characteristics and conditions. To this end, the substrate organization algorithms should be pertinent to specific management tasks/substrates in order to achieve efficient and robust operation.

The communication model used between the nodes of the intelligent substrate (horizontal interaction) or across substrates (vertical interaction) is another key aspect to be addressed. Nodes within or across substrates communicate for acquiring network awareness and for making and enforcing cooperative decisions. Current approaches in traditional management frameworks use protocol-based models and distributed object technologies, adopting mostly a client/server model. In the proposed framework, we envisage the use of alternative communication models, including aggregation and gossiping protocols, as well as the emerging information-centric publish subscribe models for the distribution of management information such as monitored data and configuration instructions. The adoption of such models will lead to a robust and efficient management information distribution mechanism where all the management components in the intelligent substrates communicate in a loosely-coupled, asynchronous manner [6]. Finally, the continuous self-organization of the substrate will benefit from such communication models since nodes will not need to be aware of the newly selected manager nodes, given that any node will be reachable using information-centric criteria.

Coordinated Decision Making

The manager nodes of an intelligent substrate need to cooperate with each other when making decisions in order to achieve a common objective with overall optimality. Each manager would normally have partial knowledge of global information and limited interaction capabilities with other managers. In this case, even "optimal" local decisions made by individual managers based on their own views of the network conditions often do not lead to global optimality. In addition, unawareness of local decisions between manager nodes may also result in suboptimal performance or even stability and convergence

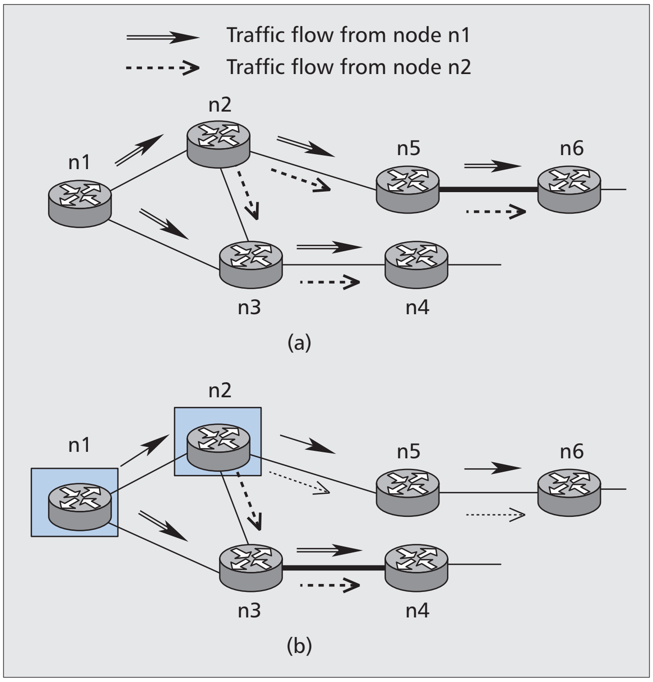

Figure 3. *Example of conflicting decisions between ingress nodes* n1 *and* n2.

issues. As such, it is required that not only necessary network condition information, but also local *intentions* by individual decision makers (i.e. managers) should be appropriately disseminated and reasoned in a distributed manner in order to reach a final coordinated network configuration outcome. We will use some examples in the next sections to illustrate how this can be achieved in specific management scenarios. The potential benefits are robustness to failures, immediate adaptability to changing conditions due to localized feedback, and automated system operation. Key research challenges are how distributed manager behavior can be coordinated and how the overall system behavior can be determined from the distributed manager interactions. Given the plethora and heterogeneity of distinct management objectives in the future advanced network management, it is not our intention to illustrate in detail any specific coordinated decision making techniques in this article.

In addition to the *horizontal* coordination between manager nodes involved in the same intelligent substrate, another key issue is the *vertical* interaction across multiple parallel substrates, each with distinct management objectives. Coexisting substrates with different but related management objectives should coordinate with each other during re-configuration. The ultimate goal is not to optimize the performance of individual intelligent substrates, but instead to achieve optimal performance and stability across all substrates by taking into account all the optimization objectives and requirements. For instance, we may consider two parallel intelligent substrates for dynamic traffic engineering and energy efficiency, respectively. The objective of the traffic engineering substrate is to perform load balancing across *all* the network links for congestion avoidance, while the strategy for energy efficiency is to provide the opportunity for traffic flows to concentrate only on a *subset* of network links, enabling the rest to enter sleep mode or to reduce their transmission rate (e.g. as proposed in [7]). Techniques such as Nash equilibrium-based game theory can be applied between individual substrates, executed by common network elements serving as manager nodes, and making sure in this case by the external management system that there is such an intersection.

Future Network Management Scenarios

Having described the features of the proposed framework, this section puts them into perspective via three representative scenarios that we believe will be typical in future networks and, more generally, the future Internet.

Adaptive Resource Management

Current resource management practices for fixed networks mainly rely on off-line approaches whereby a centralized management system computes and enforces routing configurations, based on estimated traffic demands, over long time-scales. Due to their static nature, these practices do not take network and traffic dynamics into account and can thus lead to sub-optimal network performance. To cope with unexpected traffic variations and network dynamics, approaches that can dynamically adapt routing configuration and traffic distribution are required. Existing on-line approaches have mainly focused on solutions by which deciding entities act *independently* from each other, e.g. [8]. However, this can cause configuration instabilities since decisions are based only on local information. This section describes the use of the proposed intelligent substrate for performing *coordinated* adaptive resource management in IP networks in a decentralized manner.

The deployment of this substrate aims at achieving optimum network performance in terms of resource utilization by dynamically adapting the traffic distribution according to real-time network conditions. Re-configurations occur at network ingresses (source nodes), which change the splitting ratios of traffic flows across multiple paths between source-destination (S-D) pairs. This functionality is provided by the resource management substrate, embedded in ingress nodes, which executes a re-configuration algorithm with the objective of shifting traffic from the most utilized links toward less loaded parts of the network. Performing a re-configuration involves adjusting the traffic splitting ratios of some flows for which traffic is routed across the link with the maximum utilization in the network. This results in more traffic being assigned to alternative, less loaded, paths for a S-D pair.

Due to the limited network views of individual source nodes, actions taken by more than one node at a time may lead to inconsistent configurations. For instance, in the process of shifting traffic away from highly utilized links, the different reacting nodes can re-direct traffic flows toward the same links, as depicted in Fig. 3, potentially causing new congestion. In the example, source nodes *n1* and *n2* both contribute to the load of link $l_{5\text{-}6}$, (Fig. 3a) which becomes the most utilized link in the network. If both ingress nodes react by performing re-configurations locally, more traffic will be routed toward link $l_{3\text{-}4}$ (as alternative paths to reach their original destinations), which can then become overloaded (Fig. 3b). To avoid such inconsistent decisions, only one source node is allowed to perform splitting adjustments at any time in an iterative way until no further improvement can be achieved. The resource management substrate facilitates the selection of the appropriate ingress node to perform a re-configuration of one of its local traffic flows by means of coordination messages. These communicate information within the substrate regarding the most utilized link in the network and the association of ingress nodes to that link.

The formation of the resource management substrate is based on the identification of ingress nodes in the physical network. Each node of the substrate is associated with a set of neighbors — nodes that are directly connected — with direct communication only possible between neighboring nodes. Different models can be used for the organization of the sub-

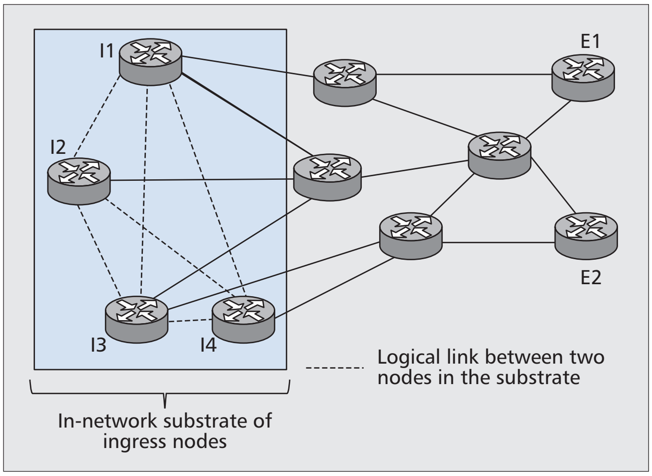

Figure 4. *Example of a network and its associated full-mesh in-network substrate of ingress nodes.*

strate, the choice of which can be driven by parameters related to the physical network, such as its topology and the number of source nodes, but also by the constraints of the coordination mechanism and the associated communication protocol. The number and frequency of messages exchanged, for example, are factors that influence the choice of model. Figure 4 gives an example of a full-mesh structure, where a direct logical link exists between the four ingress nodes (*I1–I4*) of the physical network implementing the substrate. This model offers flexibility in the choice of neighbors with which to communicate since all source nodes belong to the set of neighbors.

Energy-aware Network Management

In recent years various proposals have been made toward the realization of energy-aware ISP network infrastructures. For instance, network devices such as routers or switches can adaptively reduce their transmission rates, or even enter sleep mode upon low traffic load conditions in order to conserve energy during idle periods [7, 9]. We now illustrate how the proposed intelligent substrate paradigm can play a role in adaptively supporting optimized energy-aware operations in dynamic environments. The following constraints need to be taken into account by self-managed network elements that are able to make "green" decisions for minimizing energy consumption during operation. First, the working network topology should remain connected after some devices go to sleep, and second, the reduced network capability should not incur deteriorating service and network performances, for instance incurring traffic congestion. We first illustrate the envisioned "green functionality" in future self-managed network infrastructures through two simple examples, followed by our elaboration on how the intelligent substrate paradigm is able to support relevant operation.

In Fig. 5a, both core routers *c* and *d* have detected very low incoming traffic volume from their own upstream routers (*a* and *b*, respectively). In case both routers take the opportunity to enter sleep mode without any knowledge of each other's decisions, the working topology will become disconnected and user traffic from ingress nodes (*I1* to *I3*) will not be able to reach egress router *E1*. To avoid such a situation, the two routers need to coordinate with each other for conflict-free decision making, for instance to allow only one of them to go to sleep mode, or both routers to simultaneously reduce their transmission rates while still remaining "alive." Figure 5b illustrates another example where coordination is needed in order to improve energy-saving efficiency. Let's assume that routers *c* and *d* are currently in sleep mode. Their upstream neighbors *a* and *b* have both detected traffic spikes from ingress routers and hence both may decide to trigger their own downstream neighbors to wake up. However, in case the traffic upsurge is not sufficiently high, it might be the case that the wakeup of only one of the two downstream sleeping nodes will be able to accommodate the spike, while leaving the other in sleep mode. As such, routers *a* and *b* may also need to coordinate with each other in order to make optimal decisions for maximizing energy savings.

We now look in detail at how the intelligent substrate infrastructure is able to play a role in enabling optimized decision-making by network elements in such a distributed and dynamic environment. From the examples above, we can realize that the following information is necessary for individual devices to make coordinated green decisions: the network topology, current traffic load conditions, and the current working condition of other nodes (e.g. sleeping or active,

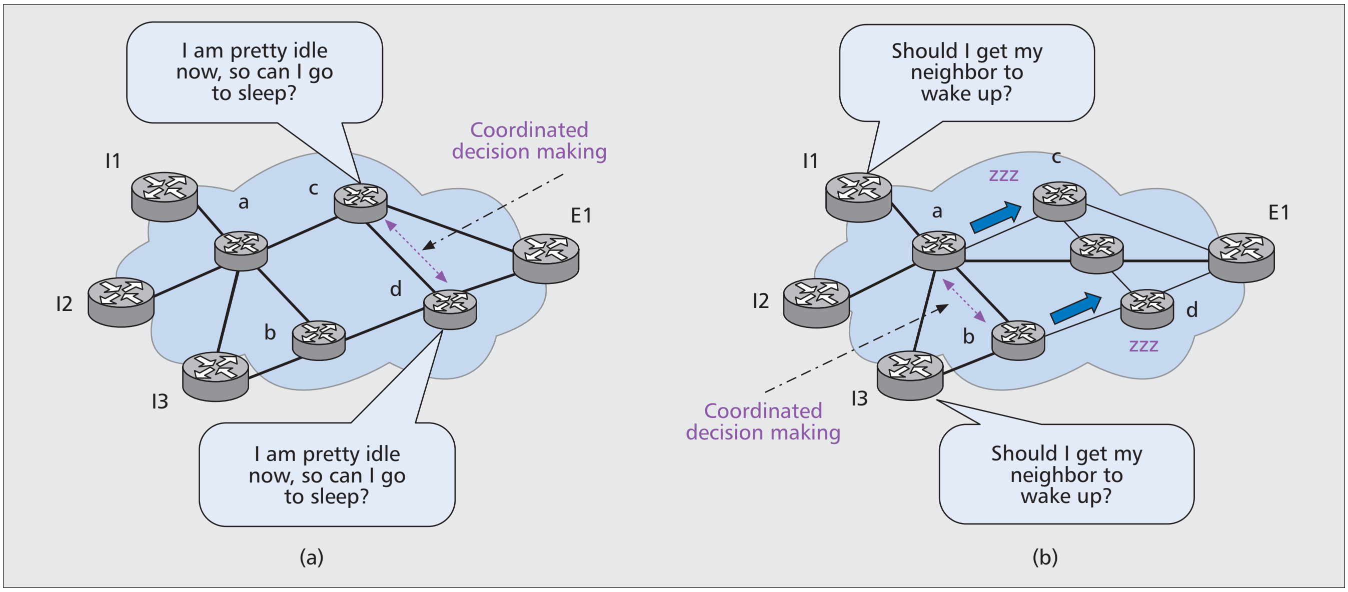

Figure 5. *Two types of router-level coordination.*

whether working with reduced transmission rates, etc.) In terms of information dissemination, up-to-date network conditions can be propagated across all routers in a time-driven manner. As an incremental solution, some existing routing protocols such as OSPF-TE, possibly with moderate extensions, are ideal vehicles for carrying such information throughout the network. Otherwise, dedicated network monitoring functions embedded in the intelligent substrate can be regarded as an alternative option for providing necessary input. Thereafter, individual routers may start to consider potential options for energy saving actions. Typically, the coordination process in the intelligent substrate follows a sequence of local decision-making process among involved elements. For instance, if a router has first detected the opportunity to perform a green operation, it needs to notify other routers about its intention. Subsequently, other routers may take into account their own local conditions in determining whether such a decision will affect the network performance. This feature requires coordinated reasoning functionality for decision reconciliation. When all the relevant routers have converged to a consistent overall decision, (some of) the devices that have made the requests will be able to activate their respective green operations.

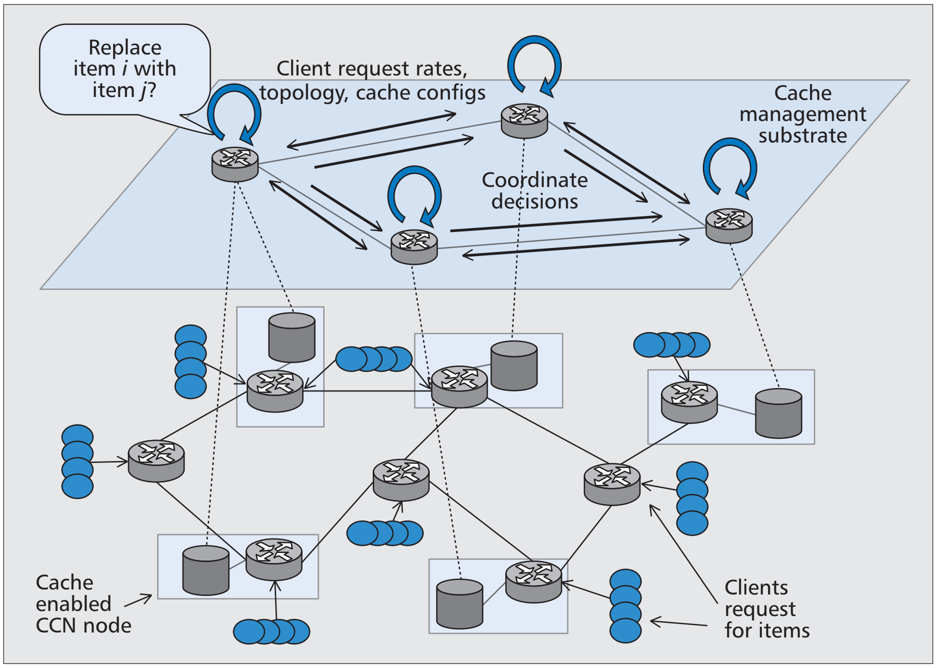

Figure 6. *Cache management substrate in content-centric networks.*

Cache Management in Content-Centric Networks

Although there is not yet consensus of what the future Internet will be, content-centric networking is emerging as the key aspect of the future networking environment given that the vast majority of activities over the Internet relate to content access and delivery. This has led researchers to point out the possibility of direct content-based routing [10], using location independent content IDs instead of node addresses. In the interim, overlay content-aware approaches are being devised for locating content instances in the case of replication and for streaming real-time content to the requesting users with appropriate quality.

In the emerging content-centric networking proposals, content is replicated almost ubiquitously throughout the network with subsequent optimal content delivery to the requesting users. Thus, efficient placement and replication of content to caches installed in network nodes is key to delivering on this promise. When a client is interested in a particular piece of content, his/her request can be redirected to one of the existing replicas rather than requiring retrieval from the original publisher. Consequently, management of such networks entails managing the placement and assignment of content in caches available in the network with objectives such as minimizing the content access latency from clients, maximizing the traffic volume served by caches, and thus minimizing bandwidth cost and network congestion.

Current approaches applied to Content Distribution Networks follow static off-line approaches with algorithms that decide the optimal location of caches and the assignment of content objects and their replicas to those caches based on predictions of content requests by users. In contrast, the deployment of the proposed intelligent substrate architecture will enable the assignment of content objects to caches to take place in real-time, based on changing user demand patterns. Distributed managers will decide the objects every cache stores by forming a substrate that can be organized either in a hierarchical manner for scalability reasons or in a peer-to-peer organizational structure. Communication of information related to request rates, popularity/locality of content objects, and current cache configurations will take place between the distributed cache managers through the intelligent substrate functionality.

Every cache manager, as depicted in Fig. 6, should decide in a coordinated manner with other managers whether to store an object that will probably lead in replacing another item already stored depending on the cache size. The decision of this swapping of stored items can be based on maximizing an overall network-wide utility function (e.g. the gain in network traffic), which means every node should calculate the gain the replacement of an object will achieve. This approach assumes that every cache manager has a holistic network-wide view of all the cache configurations and relevant request patterns, and this information should be exchanged periodically or in an event-based manner when a manager changes the configuration of its cache.

Other approaches can also be realized in which managers base their decisions on a local view of the user demand for specific objects but coordinate to maximize the overall network gain, as well as solutions where managers act selfishly aiming at maximizing their own local utility. The latter are usually formulated as strategic games [11] while an example of a distributed cache management algorithm has been proposed in [12]. Since all the above decisions are made in a distributed manner, uncoordinated decisions could lead to suboptimal and inconsistent configurations. Coordinated decision making of a distributed cache management solution can be achieved through the substrate mechanisms, by ensuring that managers change the configuration of their cache in an iterative manner until convergence to an equilibrium state is achieved.

Summary

Current practices for the configuration of ISP networks rely mainly on off-line predictive approaches, with management systems being external to the network. These are incapable of maintaining optimal configurations in the face of changing or unforeseen traffic demands and network conditions, and due to their rigidity they cannot easily support the requirements of emerging applications and future network operations. Self-management has been proposed as a potential solution to these challenges bringing intelligence into the network and thus enabling customized management tasks in a flexible and adaptive manner.

In this article we described the concept of intelligent substrates as a key component of self-management architectures since they provide the means for embedding management logic within the network. These can be viewed as a "virtualization" of management functions over the physical network infrastructure, enabling continuous self-optimization operations. We did focus on important research challenges associated with the proposed concept, including organizational and communication issues but also stability issues given both the distributed nature of the substrate and the coexistence of multiple substrates that realize different management tasks. These challenges were exemplified through representative scenarios in which we illustrated the potential role of the proposed paradigm and proposed relevant solutions. The network-wide knowledge and cooperation achieved through the intelligent substrates may significantly improve the performance of the distributed management algorithms and also reduce their execution (convergence) time. The inherent cost of additional message exchanges and computational effort can be reduced through intelligent substrate organization and communication mechanisms. We hope that the research challenges associated with the proposed paradigm will be widely addressed, and relevant in-network self-management functionality will become a reality in the medium to long term. As part of our future work we plan to investigate the pros and cons of specific organizational, communication and coordination models in the context of the three presented scenarios, identifying the required key generic infrastructure of the proposed paradigm.

References

[1] J. O. Kephart and D. M. Chess, "The Vision of Autonomic Computing," *IEEE Computer*, vol. 36, no. 1, 2003, pp. 41–50.

[2] B. Jennings *et al.*, "Towards Autonomic Management of Communications Networks," *IEEE Commun. Mag.*, vol. 45, no. 10, 2007, pp. 112–21.

[3] N. Tcholtchev, M. Grajzer, and B. Vidalenc, "Towards a Unified Architecture for Resilience, Survivability and Autonomic Fault-Management for Self-managing Networks," *Proc. ICSOC/ServiceWave Wksps.*, 2009.

[4] Y. Cheng *et al.*, "A Generic Architecture for Autonomic Service and Network Management," *Computer Communications*, vol. 29, No. 18, 2006, pp. 3691–709.

[5] A. Tizghadam and A. Leon-Garcia, "AORTA: Autonomic Network Control and Management System," *Proc. IEEE INFOCOM Wksps.*, 2008.

[6] V. Sourlas *et al.*, "Distribute, Store and Retrieve Management Policies on Wireless Ad-Hoc Networks using the Content Delivery Publish/Subscribe Paradigm," *Proc. IEEE ACNM*, 2009.

[7] M. Chang *et al.*, "GreenTE: Power-aware Traffic Engineering," *Proc. IEEE ICNP*, 2010.

[8] A. Kvalbein, C. Dovrolis, and C. Muthu, "Multipath Load-Adaptive Routing: Putting the Emphasis on Robustness and Simplicity," *Proc. IEEE ICNP*, 2009.

[9] N. Vasic and D. Kostic, "Energy-Aware Traffic Engineering," *Proc. ACM E-Energy Wksp.*, 2010.

[10] V. Jacbson *et al.*, "Networking Named Content," *Proc. ACM CoNext*, 2009.

[11] N. Laoutaris *et al.*, "Distributed Selfish Replication," *IEEE Trans. Parallel Distributed Systems*, vol. 17, Issue 12, 2006, pp. 1401–13.

[12] S. Borst, V. Gupta, and A. Walid, "Distributed Caching Algorithms for Content Distribution Networks," *Proc. IEEE INFOCOM*, 2010.

Biographies

MARINOS CHARALAMBIDES (m.charalambides@ee.ucl.ac.uk) received a BEng (First Class Hons.) in Electronic and Electrical Engineering, an MSc (Distinction) in Communications Networks and Software, and a Ph.D. in Policy-based Management, all from the University of Surrey, UK, in 2001, 2002 and 2009, respectively. He has been working as a research associate on a number of European and UK research projects since 2005, and his research interests include autonomic networking, policy analysis and refinement, and IP quality of service.

GEORGE PAVLOU (g.pavlou@ee.ucl.ac.uk) is a Professor of Communication Networks at the Department of Electronic and Electrical Engineering, University College London, UK. He holds a M.Eng. in Engineering from the National Technical University of Athens, Greece, and M.Sc. and Ph.D. degrees in Computer Science from University College London, UK. He has been responsible for a number of European and UK research projects. His research interests focus on networking, network management and service engineering. He has contributed to standardization activities in ISO, ITU-T, TMF, OMG and IETF.

PARIS FLEGKAS (pflegkas@uth.gr) received a Diploma in Electrical and Computer Engineering from Aristotle University, Greece, an M.Sc. in Telematics, and a Ph.D. from the University of Surrey in 1998, 1999, and 2005, respectively. He is currently an adjunct lecturer in the department of Computer and Communication Engineering, University of Thessaly, Greece and a postdoctoral researcher in the Centre for Research and Technology Hellas (CERTH). His research interests are in the areas of information-centric networking, policy-based network and service management, and future network architectures.

NING WANG (n.wang@surrey.ac.uk) is a lecturer at the Centre for Communication Systems Research (CCSR), University of Surrey in UK. He received his B.Eng. (Honours) degree from Changchun University of Science and Technology, P.R. China, in 1996, his M.Eng. degree from Nanyang Technological University, Singapore, in 2000, and his Ph.D. degree from the University of Surrey in 2004. His research interests include network optimization, Internet traffic engineering and content centric networks.

DAPHNE TUNCER (d.tuncer@ee.ucl.ac.uk) received a "Diplôme d'Ingénieur de Télécom SudParis" (equivalent of a Master's Degree in IT Engineering) from Télécom SudParis, FRANCE, in 2009. She is currently a Ph.D. candidate in the Communications and Information Systems Group, at the Electronic and Electrical Department, University College London, UK. Her research interests are in the areas of autonomic communications, network self-management, and on-line traffic engineering.