

Theory and Practice of Behavior in Open Computational Systems

Martin Fredriksson

Department of Software Engineering and
Computer Science, Blekinge Institute of
Technology, S-372 25 Ronneby, Sweden.

email: martin.fredriksson@ipd.bth.se

Rune Gustavsson

Department of Software Engineering and
Computer Science, Blekinge Institute of
Technology, S-372 25 Ronneby, Sweden.

email: rune.gustavsson@ipd.bth.se

Abstract

Design and maintenance of future open computational systems calls for a reassessment of current methodological approaches, theories, and practice. We have identified shortcomings in contemporary approaches, in terms of a too strong focus on exclusive models of system behavior. However, we argue that the same set of approaches also exhibits a commonality in the powerful abstraction of *domains*. In effect, we advocate the incorporation of this conception by means of a methodological focal point at all levels of behavior in open computational systems. We illustrate this perspective in practice by describing the general outline of explanatory and regulatory principles of behavior in a service-oriented layered architecture for communicating entities (SOLACE).

1 Introduction

In this paper we addressed the issue of a coherent approach towards understanding and engineering behavior in open distributed computational systems, where complexity is due to dynamics of interaction and re-configuration [Fredriksson and Gustavsson, 2001]. In literature, examples of such systems are perceived as relying on an emergent and open infrastructure of networked and embedded computers [Estrin *et al.*, 2000; Tennenhouse, 2000], as well as constituted by an evolving set of distributed, cooperating, and reusable services [Gustavsson *et al.*, 2001]. We argue that our understanding of behavior in this class of complex systems necessarily must be grounded in theory and practice of explanatory and regulatory principles. In short, we advocate a scientific-engineering approach towards information systems similar to the approaches taken by scientists-engineers in natural sciences and technologies. As such, our argument reflects a system view that typically would be characterized as a methodological approach towards observation and construction of behavior in multiagent systems. Furthermore, our approach aims at a situation where we as scientists are in possession of a set of principles that not only enable us to observe and explain behavior in some particular system. It also aims at a situation where we as engineers are in possession of a corresponding set

of principles, enabling us to actually construct behavior in this particular class of systems.

When it comes to construction or observation of open computational systems, a multiagent system approach is often considered to sufficiently capture the relevant dimensions of system behavior at the design phase. That is, modeling systems comprised by interacting agents that proliferate in a distributed, dynamical, and observable environment. However, when it comes to construction and observation of behavior in such implemented systems, we tend to focus on exclusive models of interaction, i.e., models that primarily only reflect one of possibly several perspectives of system behavior. In principle, this leads to inconsistencies in our understanding of system behavior, and even more alarming, it leads to the development of heterogeneous interaction environments. To that end, we believe that such inconsistencies originate from a situation where we neglect the fact that system behavior primarily is a natural and emergent property of system evolution, and only secondarily possible to characterize by means of abstract models at the design phase. Therefore, instead of aiming at exclusive and abstract models of behavior, we argue that theory and practice of open computational systems should be grounded in behavioral principles of natural and physical systems.

As such, this approach is a direct result from certain lessons learned in applying and integrating different models of multiagent systems. Among several multiagent systems (related to *smart houses*) we have chosen two in particular to assess principles of system integration and behavior. The goal of the first multiagent system was to reduce energy consumption by means of proactive software agents that interacted with each other over the power grid [Gustavsson, 1999]. In essence, each and every agent in that particular system resided on a networked processor, embedded in some computationally empowered device, e.g., radiators, kitchen appliances, warm water heaters, etc. By means of introducing an auction-based computational market, modeling and implementation of a very efficient and robust energy management system was successfully deployed. The goal of the second multiagent system was to support the qualitative notion of comfort in buildings [Gustavsson and Fredriksson, 2001]. In that particular context, a system of interacting agent, focused on comfort models and user preferences, was

implemented. Even though these applications illustrated the power of multiagent systems approaches towards successful construction and observation of advanced distributed systems, they also highlighted certain shortcomings. Since both applications were conceived to execute on a common and open infrastructure of networked and embedded computers, it was natural to attempt an integration of the two into one uniform system – an energy management and comfort system. However, combining models of energy saving and comfort of living rendered the underlying infrastructure incapable of supporting different analysis and design models. We now understand that this shortcoming was a natural result from an implicitly stated goal of constructing evolving, dynamical, and open multiagent systems without explicitly supporting such a primitive notion as that of *separation of concerns* by means of the underlying interaction fabric. In essence, if we aim at a common framework for theory and practice of behavior in open computational system, it is imperative that we make a clear-cut separation between actual system behavior and our temporal perception of it.

In the following sections, we propose a methodological approach that explicitly addresses conceptions of system behavior. In particular, a common framework for conceptual and physical systems is introduced in Section 2. As such, this framework introduces us to the common cognitive focal point of behavior in conceptual and physical systems, i.e., bounding spaces, and is discussed in more detail in Section 3. In order to conclude our framework of physical and conceptual bounding spaces, the focus of Section 4 is an outline of our general design and implementation of an interaction platform; developed to explicitly support the regulatory and explanatory principles of behavior in open computational systems. Finally, the material presented in this paper is discussed by means of a summary and a number of concluding remarks in Section 5.

2 Conceptions of system behavior

By means of prior experience in implementing and integrating functional components of different multiagent systems, we believe that understanding of behavior in open computational systems calls for an explicit focus on the primitive principle of separating concerns. In doing so, we argue that it is important that we understand and emphasize the common conceptions and origin of behavior in conceptual and physical systems, i.e., in order to identify a common and stable framework for design and maintenance of behavior in open computational systems. In the following section, we therefore describe our basic idea of a commonality in characterizing behavior in both physical and conceptual systems.

The scientific perspective of behavior in multiagent systems primarily involves a focus on abstract models of interaction, i.e., models that emphasize a strong focus on interactions taking place between specific system components. From such a perspective, models

are primarily applied in observation and analysis of system behavior. An issue of utmost concern is, consequently, whether or not a particular component can be classified according to the dimensions of interest, in terms of the model currently applied. However, from such a perspective the dynamical nature of a component's membership and effect on system behavior is primarily a question of conceptual adherence. That is, any *dimension* can be used to capture the notion of an agent belonging to a system and, consequently, that its behavior possibly can affect the overall behavior of the particular system in question. Therefore, as long as the agent exhibits a certain conceptual *property*, it is considered to be a physical part of a particular system. For example, an agent that interacts by means of an auction-based communication protocol could possibly be characterized as part of an auction-based computational market, but it could just as well interact with its environment by means of a completely different interaction protocol. In this sense, we consider the conceptual *domain* of a system to be of an *open* nature if agents, exhibiting a certain property, can enter or leave the conceptual system's physical boundaries. However, the idea of a fixed set of conceptual properties, identifying all the constituents forming the behavior of some particular system, is based on a very subjective notion of the dimensions that most appropriately capture our understanding of behavior in some system. Still, the interactions taking place among the constituents in any conceptual system can be characterized according to the following primitives of relevance: *domain*, *dimensions*, and *properties* [Gärdenfors, 2000]. Furthermore, we consider these primitives to form the basic structure of a model, i.e., one of possibly several conceptual perspectives of the same phenomenon.

Another example of situations where our understanding of system behavior plays an important role is in engineering of natural and physical systems. In such a context, we seek to construct systems by means of autonomous agents that through interaction form a coherent behavior. In a manner similar to that of observation and analysis, engineering of system behavior also involves a strong focus on our conception of system constituents and their possible interactions, i.e., abstract models. However, in engineering of system behavior, models are primarily applied during the design phase of construction and later on the models have fulfilled their purpose and are therefore simply discarded. Consequently, the dynamical nature of an agent's membership and possible effect on system behavior is decided upon prior to deployment. As such, a fixed set of *concepts* and their physical manifestations, or *entities*, is therefore considered to be the sole origin of system behavior. However, by discarding the models of behavior at system deployment we implicitly consider the environment, or *bounding space*, of a particular multiagent system to be of a closed nature. That is, the notion of a fixed set of conceptions, identifying the constituents that supposedly form the behavior of an open physical system, is based on a very exclusive notion of possible domains and dimensions that can be

applied in observation and analysis of system behavior.

In essence, we consider abstract models to play two distinct roles in the continuous evolution of open computational systems. On the one hand, we have the activity of observation and analysis of *a priori* existing systems and, on the other hand, we have the activity of constructing and introducing subsystems with predictable effects on their environment. In such a setting, the problem we are currently facing is very much related to the notion of *openness* as a primitive dimension of system behavior. If we are dealing with physical systems of an open nature, this means that whatever subsystem we introduce it can be both subject to and origin of behavior in the system as a whole. Furthermore, if we are dealing with conceptual systems of an open nature, this means that whatever subsystem we would like to analyze it must be possible to deduce information about the dynamical set of constituents by means of observation. In practice, the problem of theory and practice of behavior in open computational systems is twofold: (1) models of behavior are discarded prior to system deployment and (2) models of behavior are not accessible after system deployment. To that end, we argue that support for incorporation of and accessibility to models of behavior is one of the primitive roles of a *fabric of interaction*. In principle, we want to *merge* the notion of conceptual and physical bounding spaces. At this point it is therefore important to acknowledge which type of bounding space that has precedence over the other. We argue that the class of bounding spaces that necessarily has precedence over the other is *physical* bounding spaces. The reason for this is the fact that they are of a given nature, i.e., it is an indisputable and objective fact that nature exists and whether or not a conceptual construct of our minds exists is a subjectively perceived truth. This line of reasoning results in a situation where we have established that our understanding of behavior in open computational systems necessarily has to be grounded in regulatory principles of nature, i.e., the *principle of local interaction* in natural and physical systems. Furthermore, we have acknowledged that *any* model of system behavior is of a secondary order whereas the actual behavior in some system is of a primary order.

3 Primitive bounding spaces

In order to solve the problem of an inconsistency in theory and practice of understanding behavior in open computational systems we propose a *coupling of conceptions*. That is, we consider the most primitive commonality of the two approaches to be that of open bounding spaces. By means of creating an underlying fabric of interaction, governing the behavior of an open computational system, we will set the scene for our understanding of system behavior. In this section we outline our general understanding of such an interaction fabric, based upon regulatory principles of physical bounding spaces, and a corresponding cognitive system, based upon explanatory principles of conceptual bounding spaces.

Interactions are at the core of system behavior. Furthermore, from the perspective of interaction in a physical setting, we consider the behavior of a system to be primarily characterized by means of its uniqueness. That is, the uniqueness of a system is ascribed to it with respect to potential interactions with an external environment. Consequently, it is important that we are able to specify the boundaries of such an environment. If this is not possible, the concept of unique and separately existing systems will lose its meaning and our ability to model the corresponding behavior as an effect of local and external interaction will be negatively affected. Furthermore, we consider the fundamental properties shared by all *physically grounded systems* to be space and time. We should therefore primarily characterize the environment of a system to be every physical occurrence confined within the proximity of a particular bounding space. Obviously, the notions of physical grounding and bounding spaces are crucial in our identification of a uniform and coherent interaction fabric. From the perspective of an interaction fabric, the physical bounding space of some system is considered to support its basic existence and continuous activity. It is only within the proximity of such a bounding space that any *local interaction* between two system localities can take place. Consequently, it is also from the perspective of an interaction fabric that we are able to understand the notion of interaction between two distributed system localities. In practice, it would be a violation against the regulatory principles of physical systems to allow for any interaction other than that of a localized nature. In effect, if two distributed system entities can interact, this means that the underlying fabric of interaction must support the notion of *autonomous mobility*. Any system entity must be able, by means of the interaction fabric, to move from one physical bounding space to another. However, this assumption also comes with another necessary property of bounding spaces. Bounding spaces of a physical nature must support *dynamical coupling* with each other. This property has precedence over that of autonomous mobility of system entities. Furthermore, if a physical bounding space is conceived as open, this necessarily means that the entities that form the behavior in some particular system can enter or leave the system and, hence, give rise to unpredictable behavior in the overall environment. The notion of unpredictable behavior is also the main origin of the problem addressed in this paper. If an environment is considered as constituted by an unknown set of open (physical) bounding spaces, there can never be one single model that at any arbitrary point in time will capture all relevant dimensions of behavior in the environment. In other words, if we are facing the existence of a physically open environment, this means that we have to introduce support for multidimensional representation of system behavior. We return to this issue in the material presented below in terms of conceptual spaces. In summary, physical bounding spaces, as well as local interaction, require support for the following regulatory principles from

each and every bounding space in a uniform interaction fabric:

- (1) natural coupling of bounding spaces
- (2) local interaction between bounded entities
- (3) natural mobility of bounded entities

As opposed to physical bounding spaces, our conception of cognitively grounded systems focuses on the abstract representation of a certain context or situation, and such domains are considered to be comprised by a number of dimensions, i.e., system qualities and quantities. Taken all together, those distinct dimensions are considered to represent a coherent and correct description of a particular situation at hand. However, being distinct does not necessarily sum up to the notion of being exclusive. Even though the behavior of some system can be observed and characterized along certain distinct and relevant dimensions, this does not exclude the possible existence of other dimensions that could be used to represent different or more exact aspects of some system behavior in another particular context. In respect to the previous line of reasoning, engineering of a uniform, coherent, and sustainable open computational system is much more complex than traditional engineering of computational artifacts, mainly involving finite and localized state machines within closed physical bounding spaces. The particular situation at hand can be considered as a natural consequence from the fact that the fabric of interaction for open computational systems by default is of a distributed and open nature. In our aim for support of multi-dimensional representation of concerns, we primarily consider the conception of *domains*. The reason for this is that it is a common abstraction to be found in not only the notion of conceptual spaces, but also in our conception of physical bounding spaces, with a natural precedence over subjectively perceived models and corresponding dimensions of relevance. However, we argue that the cognitive notion of domains only must be related to the internal cognitive capabilities of an agent. In other words, what an agent believes to be the focal point of its conception of the external environment must never be subject to change by any external agent in the environment. In some sense, this idea corresponds to the notion of self-awareness. The second primitive of conceptual systems is that of *concepts*, i.e., a qualitative construct of our minds, bound to our conception of the surrounding environment in terms of *dimensions*. We argue that the notion of concepts can be of a localized and internal nature as well as of an external and referential nature. The third, and final, primitive of conceptual systems are the quantitative *ports*, i.e., input and output channels that represent a cognitively perceived causal relation between two *entities*. In a similar manner as that of concepts, we argue that the notion of ports can be of a localized and internal nature as well as of an external and referential nature. When two ports are related with each other, we consider the resulting cognitive construct to corre-

spond to the conceptual behavior of an ensemble of entities, i.e., the behavior of a system. In summary, a methodological focus on conceptual bounding spaces of an open nature, resulting from a situation where systems of local physical interaction is of the essence, support for certain explanatory principles from each and every bounding space is required. By means of an interaction fabric, this support corresponds to the following entity capabilities:

- (1) construct internal domains
- (2) construct domain-related internal concepts
- (3) construct concept-related internal ports
- (4) create port relations

4 The engineering problem revisited

In principle, the problem addressed in this paper is that of designing and maintaining sustainable behavior in future open computational systems. Consequently, we argue that the most important issue to address in contemporary approaches towards behavior in such systems is coherence in the supporting infrastructures. If we acknowledge that the behavior in open computational systems emerge as a result from complex interactions in open environments, this situation must not only be enforced but also provided for by the infrastructure. If systemic properties is sought for, observation and analysis, as well as construction and engineering, must be supported in a uniform manner. The most important shortcoming identified so far in our current understanding of functional requirements imposed on infrastructures for open computational systems is *support for and enforcement of model construction and observation*. It is not a shortcoming of any single approach or application, but rather an issue that emerge when two separately developed systems are to be integrated with each other. That is, integration is sought for in order to address a previously overlooked systemic property of behavior, e.g., environment comfort by means of energy management. In essence, if we aim at deploying subsystems into an existing open environment of interacting agents, the interaction fabric must:

- (1) ... be possible to conceive as a uniform and abstract interaction environment
- (2) ... enforce the construction of new and updated system models
- (3) ... support the observation and monitoring of system models

The first class of functional requirements imposed on the interaction fabric corresponds to support for the regulatory principles as previously outlined, i.e. coupling, local interaction, and mobility. The second and third class of functional requirements is enforced and




Figure 1 Regulatory mechanisms supported at the fabric level. Fabric entities act as the physical bounding space in which system entities can interact, by means of support for *interaction* and *mobility*. Fabric entities also support integration of bounding spaces, by means of *coupling*.

support for by the interaction fabric in terms of the explanatory principles that each and every entity must adhere to. In essence, the previously outlined notion of conceptual and physical bounding spaces has been implemented by means of a system of computational entities that execute on a continuously evolving infrastructure, i.e., a *service-oriented layered architecture for communicating entities* (SOLACE). The material presented below aims at an outline of that particular architecture, as well as in what way it provides for support of the mechanisms required.

5 SOLACE

As such, our methodological approach and framework emphasizes three different perspectives of behavior in open computational systems: *fabric*, *systems*, and *models*. That is, different perspectives and aspects of behavior in the same system. In principle, we consider any open computational system to rely upon support from the interaction fabric. The most primitive property of such an interaction fabric is that it continuously must support and enforce the behavior of open computational systems. As we have previously described, the regulatory principles that comes with such a perspective can most naturally be characterized as *coupling*, *local interaction*, and *mobility* (see Figure 1):

- *Coupling*. Each physical bounding space is equipped with a set of topology-dependent discovery mechanisms and a topology-independent addressing mechanism. In essence, this separation of concerns provides for the primitive abstraction layer of an open and physical bounding space, i.e., the *fabric*. As soon as a closed bounding space is discovered, it is immediately incorporated by its fabric neighbors as part of the overall interaction environment.
- *Local interaction*. Within the proximity of a physical bounding space, interaction between local entities is provided for in a topology-independent manner. However, in order to enforce the notion of local interaction between entities, telecommunication must be dealt with in a ubiquitous manner.
- *Mobility*. In a similar manner to that of bounding space coupling, the notion of mobility must be provided for by the interaction fabric. Mobility of bounded entities is provided for by means of process migration. If a local entity requires communicating with a remote entity, it must typically do so by means of local interaction




Figure 2 Explanatory mechanisms supported by the fabric level. (1) Construct domains, (2) construct concepts, (3) relate concepts, (4) construct ports, and (5) relate ports. Note that ports necessarily must be linked to each other by means of concepts in the same domain.

and, hence, the entity is transported to the remote location.

As previously described, another aspect of support for system behavior that must be provided for by the interaction fabric is the continuous creation and observation of models related to system behavior. The explanatory principles that come with such a perspective can most naturally be characterized as construction and observation of *domains*, *concepts*, and *ports* (see Figure 2):

- *Domains*. As soon as an entity is created, it must construct an initial domain, representing its own cognitive locus. This initial domain can be of a completely subjective nature, since it is internal, and simply reflects a particular agent's view of its own context.
- *Concepts and relations*. The next step in this process of model construction is twofold, either the entity creates and relates a new concept to its internal domain, by means of a dimension, or it can relate an externally observable concept to its internal domain.
- *Ports and relations*. At this point the agent has, consequently, constructed a subjectively conceived domain with certain concepts related to it by means of qualitative dimensions. The next step in this procedure is to create and name certain quantitative ports, i.e., input and output channels that represent possible context-dependent relations between the agent itself and some temporary set of external system entities.

In SOLACE, model creation is the origin and implementation of explanatory principles and a number of mechanisms are imposed on this process by the interaction fabric. On the one hand, model construction is governed by the fabric by means of which external concepts and ports an agent is allowed to construct cognitive relations with. On the other hand, model observation is also governed by the fabric, by means of which external concepts and ports an agent can observe. In essence, each entity, or agent, in the system is in control of its own conception of the environment, but it can also, by means of sharing parts of its cognitive structure, provide support for observation and explanation to other agents in the environment. By means of the interaction fabric, support for regulatory and explanatory mechanisms explicitly addresses the notion of separating concerns.

6 Summary and concluding remarks

We propose a novel methodological approach towards understanding of behavior in distributed and open computational systems. The problem we are currently facing is that there currently exists a quite large gap between understanding behavior in conceptual systems and systems of a physical nature. The framework and corresponding interaction fabric – SOLACE – outlined in this paper therefore aims at bridging this gap. The applicability of our method is, however, not only that of a common vocabulary, but it can also be used in order to compare (seemingly) diverse infrastructures for distributed, dynamical, and open computational systems, such as service-oriented systems and peer-to-peer computing [Bieber and Carpenter, 2001; Dornfest and Brickley, 2001]. Finally, by introducing a common ground for theory and practice of open computational systems, our approach is also applicable in modeling qualitative aspects of system behavior [Elrad *et al.*, 2001; Ossher and Tarr, 2001]. In essence, we argue that contemporary approaches lack a uniform and coherent grounding in conceptual and physical settings. This shortcoming is a natural result from a situation where multiple models of system behavior is addressed, without explicitly supporting such a primitive principle as that of *separation of concerns* by means of the underlying infrastructure. The results outlined in this paper are primarily related to the identification of a commonality in behavioral conceptions of physical and conceptual systems, i.e., the notion of domains. Two classes of principles and primitive mechanisms are then presented, that necessarily must be supported by the interaction fabric, constraining and enforcing system behavior. In essence, the results presented in this paper boil down to one solution in particular of the aforementioned problem: each and every model of system behavior is a subjective interpretation carried out by any agent with cognitive capabilities in an environment. The powerful and encompassing conception of domains must therefore be embedded, by means of a primitive conceptual structure, in each and every computational entity that forms the behavior in an open computational system. By means of this architecture for open computational systems, we are currently developing the first prototype systems according to the methodological approach as proposed. Examples of these systems span over a broad range of problem domains, e.g., defense systems and electronic home health care, where agent interactions are at the core of system behavior. However, the one thing they all have in common is that the expected system behavior is understood to emerge as a result from residing in an open environment of complex interactions. Future work on methodological issues, related to topics addressed in this paper, are further studied in a forthcoming book chapter [Gustavsson and Fredriksson, 2002].

References

- [Bieber and Carpenter, 2001] Guy Bieber and Jeff Carpenter. *Introduction to service-oriented programming (rev2.1)*. <http://www.openwings.org>.
- [Dornfest and Brickley, 2001] Rael Dornfest and Dan Brickley. Metadata. In Oram, A. (ed.), *Peer-to-peer: Harnessing the power of disruptive technologies*, pp. 191–202. O'Reilly.
- [Elrad *et al.*, 2001] Tzila Elrad, Robert E. Filman, and Atef Bader. Aspect-oriented programming: Introduction. In *Communications of the ACM*, vol. 44(10), pp. 28–32. ACM Press.
- [Estrin *et al.*, 2000] Deborah Estrin, Ramesh Govindan and John S. Heidemann. Embedding the Internet: Introduction. In *Communications of the ACM*, vol. 43(5), pp. 38–41, 2000. ACM Press.
- [Fredriksson and Gustavsson, 2001] Martin Fredriksson and Rune Gustavsson. A methodological perspective on engineering of agent societies. In Omicini, A., Zambonelli, F., and Tolksdorf, R. (eds.) *Engineering societies in the agents' world*, Lecture notes in artificial intelligence, vol. 2203, 2001. Springer Verlag.
- [Gustavsson, 1999] Rune Gustavsson. Agents with power. In *Communications of the ACM*, vol. 42(3), pp. 41–47, 1999. ACM Press.
- [Gustavsson and Fredriksson, 2001] Rune Gustavsson and Martin Fredriksson. Coordination and control of computational ecosystems: A vision of the future. In Omicini, A., Zambonelli, F., Klusch, M., and Tolksdorf, R. (eds.) *Coordination of Internet agents: Models, technologies, and applications*, pp. 443–469, 2001. Springer Verlag.
- [Gustavsson and Fredriksson, 2002] Rune Gustavsson and Martin Fredriksson. Humans and complex systems: Sustainable information societies. In Olsson, M. O. and Sjöstedt, G. (eds.) *Revealing complex structures: Challenges for Swedish systems analysis*. FORMAS.
- [Gustavsson *et al.*, 2001] Rune Gustavsson, Martin Fredriksson, and Christer Rindebäck. Computational ecosystems in home health care. In Dellarocas, C., and Conte, R. (eds.) *Social order in multiagent systems, Multiagent systems, artificial societies, and simulated organizations*, vol. 2, pp. 201–220. Kluwer Academic Publishers.
- [Gärdenfors, 2000] Peter Gärdenfors. Dimensions. In *Conceptual spaces: The geometry of thought*, pp. 1–32, 2000. The MIT Press.
- [Ossher and Tarr, 2001] Harold Ossher and Perri Tarr. Using multidimensional separation of concerns to (re) shape evolving software. In *Communications of the ACM*, vol. 44(10), pp. 43–50. ACM Press.
- [Tennenhouse, 2000] David Tennenhouse. Proactive computing. In *Communications of the ACM*, vol. 43(5), pp. 43–50, 2000. ACM Press.