

# Data Access in Distributed Simulations of Multi-agent Systems\*

Dan CHEN<sup>1,2</sup>, Roland EWALD<sup>2,4</sup>, Georgios K. THEODOROPOULOS<sup>2</sup>, Robert MINSON<sup>2</sup>, Ton OGUARA<sup>2</sup>, Michael LEES<sup>3</sup>, Brian LOGAN<sup>3</sup>, and Adelinde M. UHRMACHER<sup>4</sup>

<sup>1</sup>Institute of Electrical Engineering, Yanshan University, Qinhuangdao, China

<sup>2</sup>School of Computer Science, University of Birmingham, UK

<sup>3</sup>School of Computer Science and IT, University of Nottingham, UK

<sup>4</sup>Department of Computer Science, University of Rostock, Germany

**Abstract**— Distributed simulation has emerged as an important instrument for studying large-scale complex systems. Such systems inherently consist of a large number of components, which operate in a large shared state space interacting with it in highly dynamic and unpredictable ways. Optimising access to the shared state space is crucial for achieving efficient simulation executions. Data accesses may take two forms: locating data according to a set of attribute value ranges (Range query) or locating a particular state variable from the given identifier (ID query and update). This paper proposes two alternative routing approaches, namely the address-based approach, which locates data according to their address information, and the range-based approach, whose operation is based on looking up attribute value range information along the paths to the destinations. The two algorithms are discussed and analysed in the context of PDES-MAS, a framework for the distributed simulation of multi-agent systems, which uses a hierarchical infrastructure to manage the shared state space. The paper introduces a generic meta-simulation framework which is used to perform a quantitative comparative analysis of the proposed algorithms under various circumstances.

**Index Terms**— Multi Agent Systems, Complex Systems, Distributed Simulation, Data Management, Range Query

## I. INTRODUCTION

THE last decade has witnessed an explosion of interest in complex systems, which involve dynamic and unpredictable interactions between large numbers of components including software, hardware devices (such as sensors), and social entities (people or collective bodies). Examples of such systems range from traditional embedded systems, to systems controlling critical infrastructures, such as defence, energy, health, transport and telecommunications, to biological systems, to business applications with decision-making capabilities, to social systems and services, such as e-government, e-learning etc. The complexity of such systems renders simulation modelling the only viable method to study their properties and analyse their emergent behaviour. Multi-agent systems (MAS) have emerged as a particularly suitable paradigm for modelling complex systems. When embedded in a real system, a MAS is itself a complex system whose properties

\* Contact Author: Dan Chen, Email: d.chen@cs.bham.ac.uk; chendan@pmail.ntu.edu.sg. Dan Chen and Roland Ewald had a Postdoctoral Research Fellowship and an Internship with Birmingham respectively while this work was undertaken.

and emergent behaviour have also to be analysed via simulation [7][14].

The application of agent-based simulation to ever more complex problems has placed it in the highly computation intensive world with computational requirements far exceeding the memory and performance capabilities of conventional sequential computer systems. As a result, parallel and distributed simulation emerges as a particularly promising and viable approach to alleviate the simulation bottleneck in the design and analysis of large, complex, agent-based systems.

Amongst the most influential of these approaches, the Logical Process Paradigm seeks to divide the simulation model into a network of concurrent Logical Processes (LPs), each of which models some object(s) or process(es) in the simulated system. Each LP maintains and processes a portion of the state space of the system and state changes are modelled as time-stamped events in the simulation [10].

In conventional distributed simulations, the shared state is typically small and the processes interact with each other in a small number of well-defined ways. The topology of the simulation is determined by the topology of the simulated system and its decomposition into processes, and is largely static. However, in the case of agent-based systems, which operate in complex environments and interact with it in highly dynamic and unpredictable patterns, it is often difficult to determine an appropriate simulation topology a priori. In such systems there is a very large set of shared state variables which could, in principle, be accessed or updated by the processes in the model. Encapsulating the shared state in a single process (i.e., via some centralised scheme) introduces a bottleneck, while distributing it all across the LPs (decentralised, event driven scheme) will typically result in frequent all-to-all communication and broadcasting.

In [21] we have proposed an approach to manage the shared data in distributed simulations of multi-agent systems (MAS). The approach is based on the notion of *Spheres of Influence (SoI)*

and uses a hierarchical simulation infrastructure to dynamically decompose and distribute the shared state. This framework has been realised in the context of PDES-MAS<sup>2</sup>, a system for the distributed simulation of multi-agent systems. Management of shared data in distributed simulations needs to address two problems, namely *data distribution* and *data accessing*. In [27] we addressed the first problem and described data distribution algorithms for the PDES-MAS framework; in [18][19] we discussed synchronisation issues that arise from such distributions.

In this paper, we focus on the second problem of data access. Data accesses target both individual data items (referred to as *ID queries*) and selected data items overlapping given query windows (referred to as *Range queries*). This is a challenging issue, particularly when both the value and the physical distribution of data items are dynamic. The physical distribution of data items refers to the assignment of data items to LPs, which are running in separate threads and could thus be distributed physically over a set of machines. Value distribution, in contrast, characterises the content of data items.

Typically two basic criteria to locate data items can be identified, namely the physical location of individual data items and their attribute value. This paper describes two new candidate algorithms for data accessing in the context of the PDES-MAS framework, the *address-based* and the *range-based routing*, each relying mainly on one of the two aforementioned respective criteria. Data access approaches considerably influence the efficiency of simulation execution and contribute to the complexity of system design. This paper first gives a qualitative comparison of the proposed algorithms and then provides a quantitative analysis of the dynamics of the two solutions. The performance of the candidate algorithms has been analysed with respect to the behaviour of the simulated system as well as the characteristics of the simulation infrastructure. Emphasis has been given to the problem of dynamic range queries. The experimental framework

<sup>2</sup> <http://www.cs.bham.ac.uk/research/pdesmas>

used can provide a substantial reference for system designers to address similar problems.

The two algorithms were first outlined in [41] where an initial evaluation was also presented. This paper presents a more detailed description of the algorithms and an extended set of results with an in-depth analysis. The rest of the paper is organized as follows: the PDES-MAS framework is introduced in Section II. Section III discusses the alternative approaches to data accessing. Section IV provides a qualitative comparison of the proposed routing approaches. Section V describes the meta-simulation model which has been developed to study the dynamics of the routing approaches. Section VI presents the benchmark experiments and results. Related work is briefed in Section VII. Section VIII concludes the paper with a summary and ideas for future work.

## II. THE PDES-MAS FRAMEWORK

PDES-MAS adopts a standard discrete event simulation approach with optimistic synchronization [21]. When constructing multi-agent systems using the framework, an MAS is modelled as a network of LPs (Fig. 1). In particular, each agent is modelled as an *Agent Logical Process (ALP)*. An ALP has both private state and shared state. The private state is maintained within the ALP, while the shared state can be accessed (read or updated) by other ALPs in the model. Changes in the state of an ALP that may have a causal impact on other ALPs are referred to as external events and are represented by a time-stamped operation on the shared state.

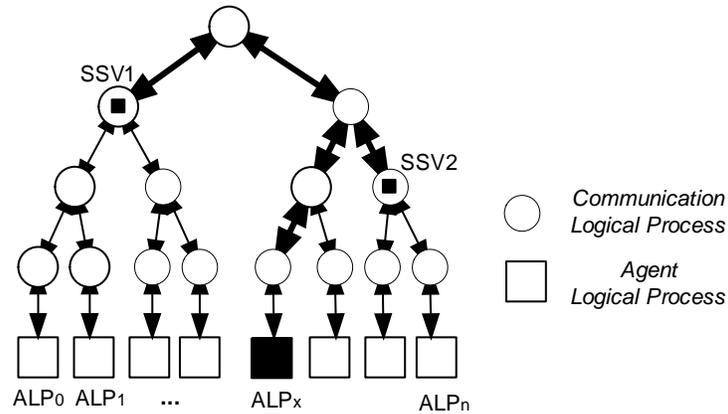


Fig. 1. Overview of the PDES-MAS framework

The shared state is modelled as a set of *Shared State Variables (SSVs)*, each of which is a tuple of the form  $\langle SSV\ ID, attribute\ type, value, timestamp \rangle$ . In PDES-MAS, the shared state is maintained by a tree-structured set of additional logical processes, *Communication Logical Processes (CLP)*, which cluster agent models and shared state according to the agents' SoIs [21]. As the access patterns on the shared state change, so does the configuration of the tree and the distribution of state (i.e., its allocation to CLPs) to reflect the logical topology of the model.

Redistribution of shared state can be achieved in a number of ways, such as by creating/deleting CLPs, by migrating ALPs through the tree, or by migrating state between CLPs. For the purposes of this paper, we choose to use a fixed tree of CLPs and move SSVs through the tree to achieve redistribution. SSVs are constantly moved closer to the ALPs that access them most frequently, reducing the total access cost and thus contributing to the scalability of the framework [27]. The framework does not make use of proxies and only one instance of an SSV is present in the tree at any particular moment. ALPs always link to the leaf CLP nodes in the tree.

The CLP tree provides common services to the ALPs, which includes: (a) facilitating the construction of the distributed simulation; (b) clustering and interoperating the ALPs; (c) managing shared data and balancing load incurred by accessing the shared state; and (d) facilitating syn-

chronization of the ALPs.

Fig. 2 illustrates the relationship between an ALP and the CLP tree. The operation of the CLP tree remains transparent to the ALPs during the simulation. The PDES-MAS framework provides a software library to the ALPs to interact with the CLP tree through two interface modules, referred to as *SimulationAmbassador* and *AgentAmbassador*. An ALP issues requests to access shared state variables through the *SimulationAmbassador* module which forwards the requests to the parent (or the server) CLP. If the required SSV is not held locally, the server CLP passes the request to its parent CLP to deal with the request. The return data and control messages (i.e., rollback) are conveyed to the ALP via its *AgentAmbassador* module.

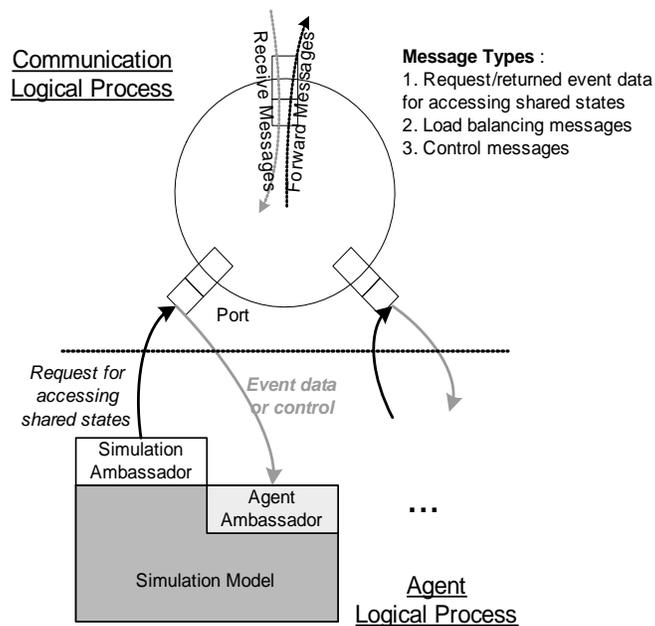


Fig. 2. Relationship between the CLP Tree (Hierarchical Simulation Infrastructure) and ALPs

Fig. 3 gives a schematic view of a CLP, which interacts with other LPs in the system via ports. Ports link the individual LPs together to form the overall PDES-MAS simulation system. In this paper, the *incoming* port of a CLP is referred to as the one from which a request on accessing SSV is received. Each CLP is also a router responsible for forwarding access requests to the destination CLP(s) that host the target data, and forwarding is via the *outgoing* ports.

The port is specially designed to maintain the distribution of the values of SSVs in the value space<sup>3</sup> classified by the types of SSVs. The *attribute value range* denotes a certain range of the values of a set of SSVs associated with a particular attribute (or a set of attributes). The “extent” covered by the attribute value range may also vary with different routing algorithms. The distribution concerns SSVs in the local CLP and/or SSVs in remote CLPs, for instance the overall system beyond a port or only the direct neighbours (parent and/or children nodes if any) to the CLP through a port (i.e., as in Figures 5 and 8).

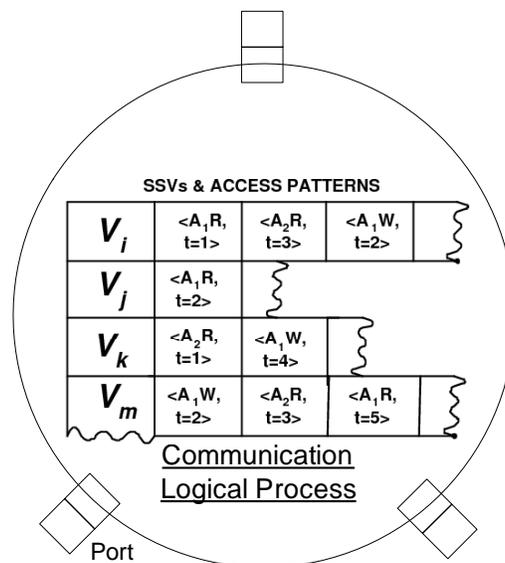


Fig. 3. Communication Logical Process and Ports

The distribution of the values gives a panorama of the status of SSVs in the system. No matter what resolution and extent are chosen for an attribute value range, it should always correctly reflect the status of SSVs and can be refreshed once the status changes. More details of attribute value range are available in Section III.

### III. DATA ACCESS IN PDES-MAS

Routing approaches are needed for ALPs and CLPs to locate (a) SSVs according to the attribute value ranges (range query) and/or (b) a particular SSV from the given ID (ID query and up-

<sup>3</sup> For example, we define “x-position” in an extent [0, 100]. Given 100 SSVs of x-position, the values of these SSVs may distribute evenly

date). The two types of locating differ from each other significantly. A range query searches for a group of SSVs based on the specified common attributes and constraints, and the targets tend to differ in each query. This is similar to multicast on dynamic groups. In contrast, ID query aims to locate a unique SSV given its ID, and therefore is closer to a unicast operation.

Routing access requests to SSVs can be performed via either their location information in the tree or the attribute value range information maintained at the ports of the CLPs. SSVs may be moved between CLPs, but there are no multiple copies of a single SSV in the overall system. The status of an SSV can be altered by updating and load management [27]. An update may change the value of the SSV, which directly affects the corresponding attribute value range. Load management may induce the migration of the SSV to a different location in the CLP tree, thus, changing the value ranges at related ports. This immediately influences ID queries on this SSV and possibly range queries.

To facilitate the location of SSVs in the CLP tree, it is helpful to encode<sup>4</sup> the tree to identify the CLPs. The address of an SSV is defined as the code of the CLP at which it is maintained (the CLP is referred to as the *host* CLP of this SSV, while the SSV is considered *local* to this CLP). When forwarding access requests, a CLP decides through which port to push the access request to the destination CLP.

The fixed architecture of a CLP tree determines that: (a) an SSV can only migrate from a CLP to its direct neighbours, and (b) between any ALP and CLP, there exists only one unique path. Once the target SSVs are located, the returned values need to be simply propagated along the path which the query just traversed in the reverse direction to the source ALP. This section presents two candidate approaches to route ID and Range queries through the tree of CLPs. The two

---

from 0 to 99, such as (0, 1, ..., 99), or concentrate on [50, 51].

<sup>4</sup>No particular coding scheme is preferred as long as it identifies each CLP uniquely and remains consistent in the simulation.

approaches dynamically adapt to different properties of the shared state and the system.

To store the information efficiently, the overall value range of each SSV type is divided into a number of *segments*. Hence, only one bit per segment is needed to store information about the existence of SSVs with values covered by this segment. For example in the case of a CLP containing a set of SSVs with values listed as: {20, 53, 56, 70, 80, 190, 310, 370} (see Fig. 13), instead of using a simple range description, such as [Min(20), Max(370)], the value space can be segmented as Seg1: [0, 100], Seg2: [100, 200], Seg3: [200, 300], Seg4: [300, 400], Seg5: [400, 500], ....etc. The approach logs the number of SSVs whose values fall onto each segment. For Fig. 4(A), the following segments are defined: {Seg1(5), Seg2(1), Seg3(0), Seg4(2), Seg5(0)}; the attribute range is {Seg1 $\cup$ Seg2 $\cup$ Seg4}. When an update occurs, even if the value of the updated SSV is beyond the original [Min, Max], the segmented range may not need to change. For instance, if the SSV with value 370 is updated to 380, the updated value still falls onto Seg4 (Fig. 4(B)). The range has to be updated only when there is no SSV having a value in one of the segments (narrowing, Fig. 4(C)) or any SSV's value exceeds all existing segments (expanding, Fig. 4(D)).

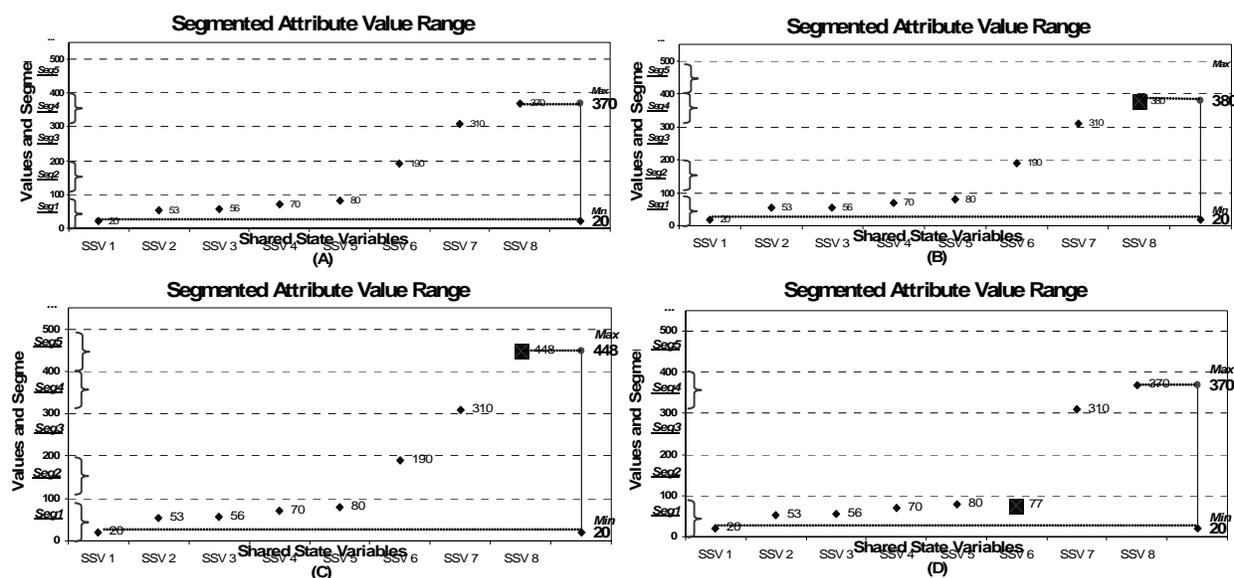


Fig. 4. Segmenting Attribute Value Range

### A. Address-based Routing

The address-based routing searches for SSVs according to their addresses, namely their exact location (host CLP) in the tree. Fig. 5 illustrates an address-based routing approach, which binds the ID of an SSV to its address. Each server CLP maintains a routing table containing the addresses of SSVs that have been accessed in the past. The routing table has a hierarchical format using attribute IDs as indices. From a particular object attribute's perspective, the table maintains the addresses of CLPs hosting the same type of SSV. The SSV IDs of this type are recorded under the host CLP entry. Furthermore, each CLP stores information about the values of SSVs that are hosted by its immediate neighbours. This information is obtained and refreshed when updates on those SSVs occur.

#### 1) Range Query with Address-based Routing

When an ALP issues a Range query (see example in Fig. 5), its server CLP propagates the request to all CLPs which host SSVs of that SSV type (in this example,  $CLP_1$  and  $CLP_2$ ). If the values of its SSVs are not within the segments covered by the Range query, the neighbours of a CLP can stop the query (unless there is another CLP that needs to be reached).

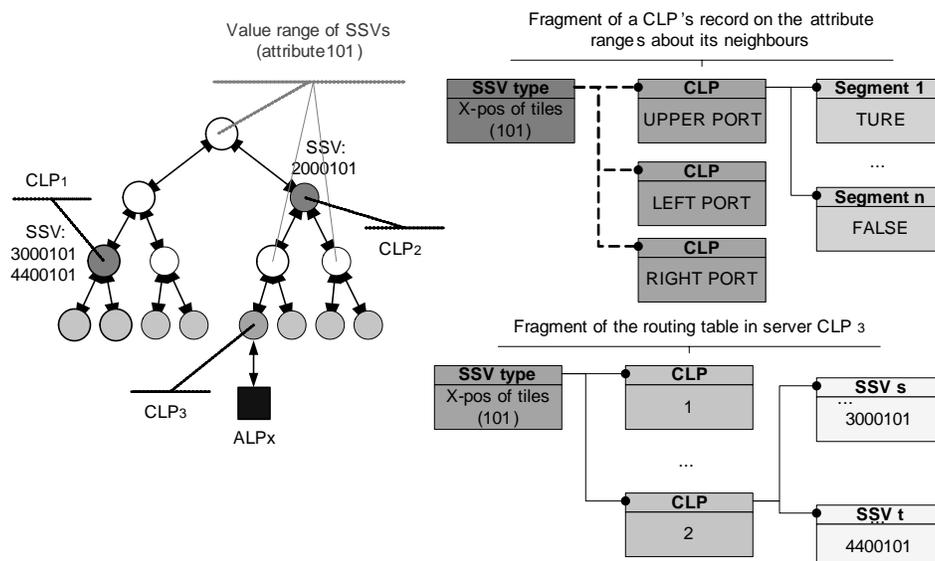


Fig. 5. Address-based Routing

The address-based routing algorithm is illustrated in Fig. 6. When a server CLP receives a range query request for accessing SSVs of a particular type, it needs to resolve the addresses of these SSVs from its routing table. If the attribute ID associated to the SSVs cannot be found in the table, the server CLP will initiate a global search for the SSVs with the designated type in the CLP tree. The routing table will then be updated with the returned results, and the current range query finishes. If the attribute ID can be found in the server CLP's routing table, the range query request will then be forwarded to the destinations. When the request reaches the CLP next to a destination, the CLP checks the query's range against the attribute range information about the destination to see whether both ranges overlap. If they don't overlap, search on the destination CLP will be given up. If they overlap, the query will be delivered to the destination CLP, and the results will be returned to the server CLP. Eventually, the server CLP consolidates results (if any) collected from itself and the other destination CLPs, and conveys those to the corresponding ALP, and the range query completes.

## *2) ID Query with Address-based Routing*

The algorithm for an ID query is straightforward. When an ALP issues an ID query, the server CLP locates the destination (if the SSV is not maintained locally) from its routing table and forwards the query to the particular host CLP.

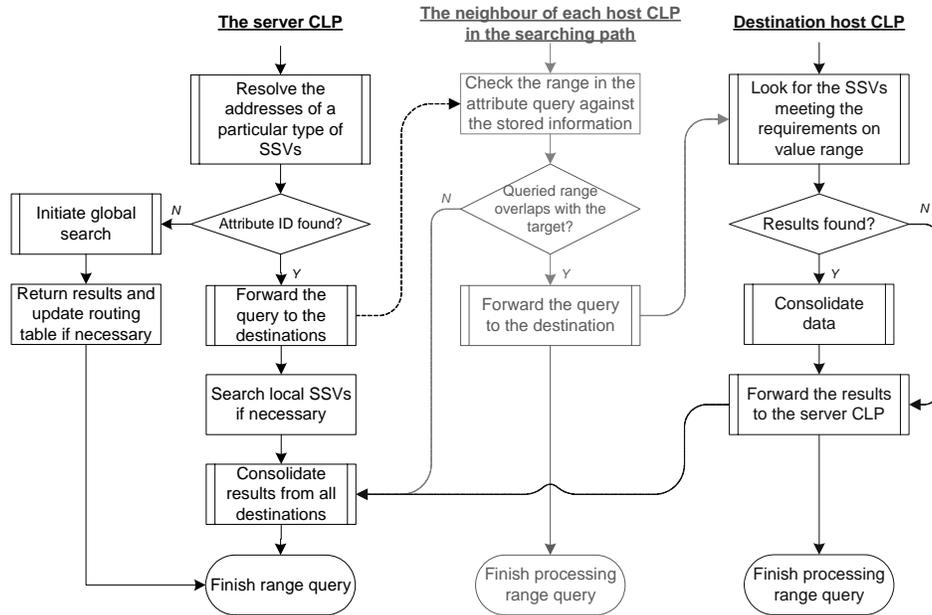


Fig. 6. An Address-based Routing Algorithm

### 3) Migration of Shared State Variables

When an SSV migrates, the change of the SSV's address immediately affects the ID-to-address binding. The routing tables in the server CLPs have to be updated with the new address. A gradual address updating scheme is used to avoid global propagation for updating addresses. This is illustrated in Fig. 7. The port through which an SSV migrates is recorded in the original host CLP, and the CLP becomes the SSV's *correspondence CLP*. The map between port and migrated SSVs is looked up as another routing table for searching those SSVs. When  $ALP_x$  originally accessed  $SSV_1$ ,  $CLP_m$  was  $SSV_1$ 's original host CLP. Obviously, at that time the query from  $ALP_x$  on  $SSV_1$  must be routed to  $CLP_m$ 's right port along a fixed path (the *original path*).

After  $SSV_1$  migrates to a new host, from  $ALP_x$ 's point of view, there are three different cases: (1)  $SSV_1$  has been pushed further away (Fig. 7(A)), (2)  $SSV_1$  has been brought closer to  $ALP_x$  (Fig. 7(B)) or (3)  $SSV_1$  has been moved elsewhere (Fig. 7(C)). For case (1), when a new query reaches  $CLP_m$ , it will be forwarded to  $CLP_m$ 's direct neighbour beyond its upper port. The forwarding will be relayed until the new host is found. For case (2), the new host must be an inter-

mediate node in the original path, and the query will be answered straightforwardly when reaching the host. For case (3), along the original path, the query will pass a correspondence  $CLP_n$ , and then it will be relayed downwards to the new host  $CLP$  of  $SSV_1$ . From the above discussion, it is clear that no matter in which new host an  $SSV$  locates, the query will only travel along the unique path between the  $ALP$  and the new host. The routing approach does not concern any other off-path  $CLP$ .

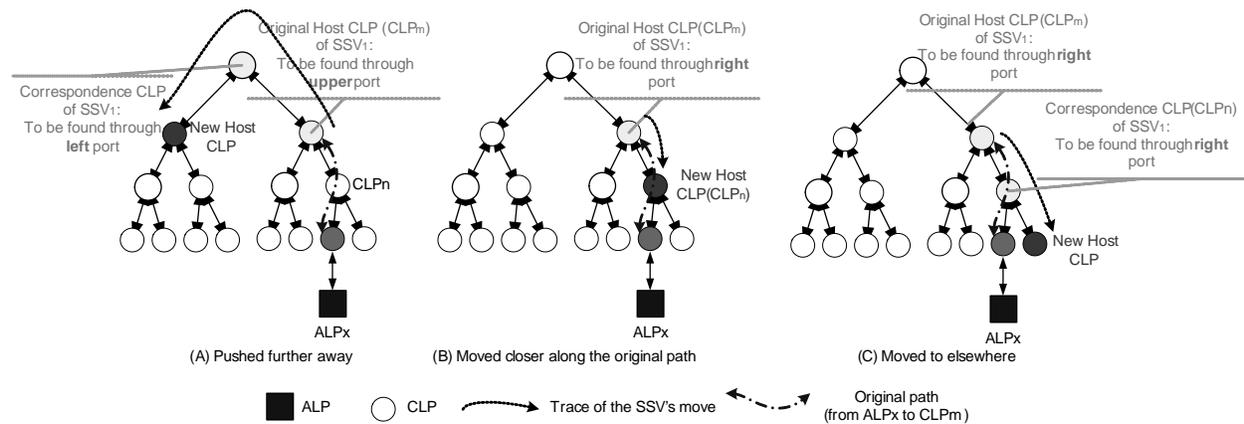


Fig. 7. Gradual Address Updating and Routing

As a result of migrating an  $SSV$ , the value range of its type may change in both the original and the new host. Therefore, the attribute value ranges related to the two hosts have to be re-computed and updated. However, this does not influence the routing on other  $SSVs$ .

### B. Range-based Routing

The range-based routing approach uses information about the attribute value range to locate  $SSVs$  in the tree. Under this approach, a  $CLP$  forwards a query according to (a) the availability of the  $SSV$  type being queried beyond its ports and (b) the value range of  $SSV(s)$  belonging to the given type.

#### 1) Range Query with Range-based Routing

The approach matches the query window with the attribute value ranges along the searching

paths to gradually approach the potential targets. When an ALP issues an attribute value range query, a range-based routing will start from its server CLP. Searching will stop at the directions where the query window and attribute value ranges do not overlap. Like the address-based routing, the range-based routing stores the value range for each port by segmenting it. But instead of only considering the neighbour CLPs, each bit marks the existence of matching SSVs somewhere beyond the port. The port information will be kept up to date according to returned messages from neighbours; if an empty message is returned, there are no matching SSVs behind the corresponding port and this information will be used in the future. Obviously, the port information may need to be updated when the value of any SSV changes.

In the example shown in Fig. 8,  $CLP_m$  keeps a record of the SSVs within the enclosed areas in its three ports respectively. Suppose that a range query for SSV type “X-pos” and range [2, 6] reaches  $CLP_m$ : The CLP has two possible directions to relay the query. Direction  $B$  will be given up as the attribute value range [8, 9] does not overlap the window [2, 6]. The query will be forwarded only towards direction  $A$ , as the corresponding attribute value range [1, 5] matches the condition. Fig. 9 depicts the range-based routing algorithm, it applies to any intermediate CLP which receives a range query request from a port. The CLP checks whether there is any SSV matching the queried value range. After that, the attribute ranges maintained by the other two ports will be compared with the queried range, the query will then be sent to the other CLPs through the ports whose attribute values overlap the range query. The CLP waits for the results (if any) returned from the other ports and conveys them together with local targets (if any) to the port via which the range query arrives. The range query operation on this CLP completes.

## 2) Shared State Variable (ID) Query with Range-based Routing

The above approach can also be applied to access SSVs by ID. In this case, the ID number range is segmented as well, so that ID queries can be resolved in a similar way as Range queries.

Furthermore, load management has direct influence on the ID range. After an SSV migrates, the ID range of its original and new host CLP may change. This issue is similar to dealing with the SSV update in attribute value range query.

#### IV. A QUALITATIVE COMPARISON

This section discusses the respective advantages and disadvantages of the two proposed approaches from a qualitative point of view.

**Availability of Shared State Variables.** Both approaches provide correct information of where SSVs may be available and avoid routing accesses to those CLPs which do not contain any SSV of the requested type. The major difference is that the address-based approach gives exact locations while the range-based approach directs the accesses to the correct searching paths.

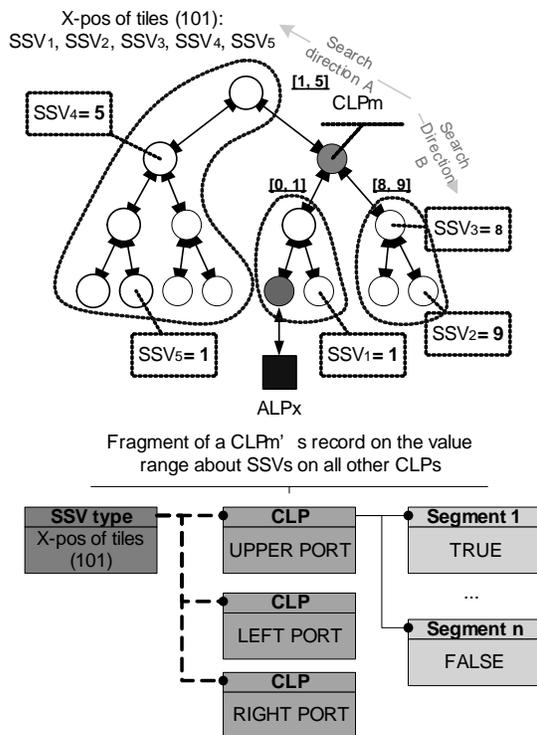


Fig. 8. Illustrating a design for Range-based Routing

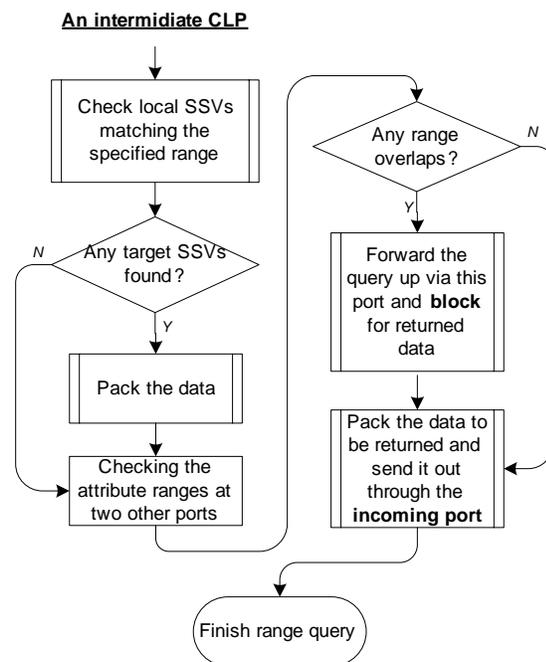


Fig. 9. A Range-based Routing Algorithm

**Efficiency in Performing Range Queries.** When performing a range query, the range-based

approach forwards the query to the potential targets in a bottom up fashion. In the target narrowing procedure, those out-of-range SSVs can be filtered out effectively. The address-based approach needs to route the access to the neighbours (in the searching path) of all CLPs to match the query and the host CLPs ranges. In general, the ranged-based approach forwards a range query to a smaller set of CLPs than the address-based solution does. The two approaches require searching within the same number of potential host CLPs. They do not differ in the complexity of searching within those hosts or the overhead in delivering results to the requesters.

**Complexity for Maintaining Range Information.** The range-based approach relies on the attribute value range information. From a CLP's point of view, the attribute value ranges (in the other CLPs beyond each of its ports) must be available and accurate. A range-based algorithm needs to manage the segmented ranges properly. A fine-grained segmentation can obtain accurate routing, while a coarser segmentation will reduce broadcasting of updated ranges. In the address-based approach, no range information is maintained or broadcasted. A CLP needs to simply compute its local attribute value range prior to notifying its direct neighbours. In the case of handling SSV migration, the address-based approach does not incur any extra communication overhead for routing. In contrast, the range-based approach has to consider the immediate impact on the previous and current host CLPs, which may involve updating attribute value ranges on the CLPs.

**Efficiency of ID Queries.** The address-based approach is able to resolve the address at the server CLP immediately for an SSV ID query, and the SSV can be accessed via a fixed path without routing to irrelevant CLPs. Using the range-based approach, querying an individual SSV is not straightforward.

**Complexity for Maintaining Routing Information.** The address-based approach assigns dif-

ferent tasks to different CLPs (servers and intermediate nodes in the tree) with server CLPs keeping addresses of all the SSVs in its client ALPs' interest. However, address resolving within a centralised node can be optimised. The address change of any SSV does not affect routing. On the other hand, the range-based approach distributes the routing information throughout the CLP tree in an implicit manner. The address changing of any SSV may affect multiple CLPs or even the entire CLP tree.

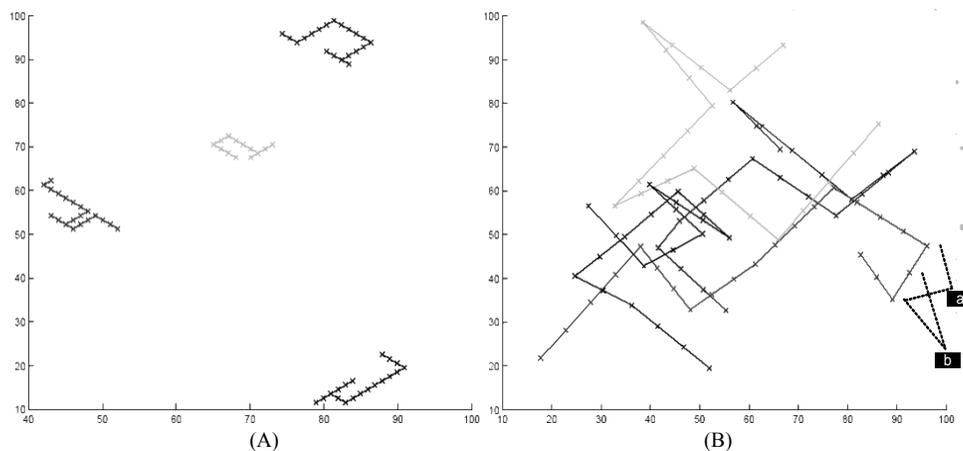


Fig. 10. Illustrating Different Agent Movement Patterns

## V. MODEL OF THE SIMULATION SYSTEM

A comparative and quantitative analysis of the two proposed approaches is a non-trivial task, as it involves the evaluation of efficiency in performing Range queries and ID queries, the complexity of maintaining routing information, and maintaining range information, design complexity etc. From the scale of the CLP tree and the number of SSVs, it is relatively straightforward to estimate the computational and communicational complexity of the address-based routing approach using mathematical approaches [7]. However, the evaluation of range-based routing needs to consider other complicated factors at both application level and simulation level.

For a quantitative analysis, one approach would be to directly implement and integrate the two approaches into the PDES-MAS kernel. However, this would require considerable implementation efforts, and at least part of the implementation could be in vain, as the strategies may not

meet the performance requirements. To avoid this we have adopted a meta-simulation approach, as proposed for instance in [16][29][30].

The design of the meta-simulation [8] follows a layered approach similar to [12]. At the top layer, we find the application model, which is responsible for generating realistic query patterns. The next layer is the middleware layer, where the routing approaches are described and the PDES-MAS framework is represented. The third layer, which typically is reserved for the model of the underlying infrastructure, is implicitly represented in the performance measurements which are integrated by calculating the costs of queries in the second layer. Thus, similar to many simulations of P2P systems, the characteristics of the underlying network are abstracted away by only counting hops and messages.

#### A. *Application Model*

The application model focuses on the simulation of situated agents, wherein an agent has a position within the model that determines its region of interest: only objects situated in the region can be accessed by the agent. In addition, situated agents are usually able to change their own positions. This behaviour was modelled for a two dimensional environment, as shown in Fig. 10. An agent moves step-wise towards a pre-selected target along the shortest path, and it randomly chooses a new target on arrival. The distance of the new target, the *target distance* (mark “a”), is defined by the number of steps it takes the agent to reach it. The distance an agent can move in each step is referred to as the *step size* (mark “b”), which reflects the rate of change of the agent’s access pattern. The step size and target distance determine the activity scope and movement speed of an agent. After each step of movement, an agent generates ID or Range queries concerning its current region of interest.

We assume that all SSV types within the MAS model have a spatial meaning, i.e. the value

ranges for Range queries reflect the actual positions of the agents. Each SSV type represents a certain dimension of the environment, such as “X-pos” or “Y-pos”. Note that this assumption should not affect the generality of our model. We assume that other SSV types, such as non-spatial attributes of the modelled objects, are accessed on demand *after* range queries identified all objects in the agent’s region of interest. SSVs may have a uniform value distribution or multiple normal distributions, illustrated in Fig. 11(A) and (B) respectively.

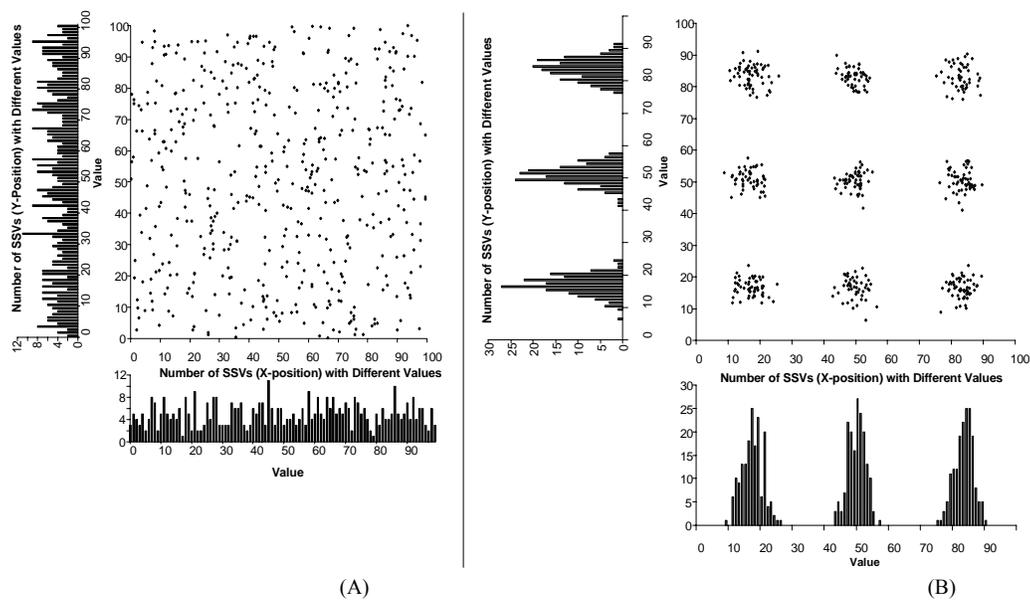


Fig. 11. Illustrating Different SSV Value Distribution Patterns

### B. PDES-MAS Model

The model of the PDES-MAS framework has been simulated using discrete time steps, and it is formed by a set of SSVs. Each SSV consists of its (unique) ID, type, value and position in the CLP tree. The modelled CLP tree is binary and complete; therefore its structure can be defined by its depth. Another important parameter is the number of segments used by both routing algorithms, which determines the granularity of the description of the value distribution of SSVs. To precisely identify the effect of each individual “impact factor”, the model adopts different SSV

distribution patterns for different runs<sup>5</sup>, while assuming that the locations of SSVs are fixed (as proposed in [27]) in each run. This avoids bias and additional parameters to be explored, since the dynamics of load balancing might be very complex. Instead, parametrisable properties of the SSV distribution are introduced in the next section. As the SSV distribution would be controlled by any load management scheme, its impact on routing algorithm performance allows to investigate the relationship between routing and load balancing in a more general way. Nevertheless, the actual performance of both routing algorithms in combination with concrete load management schemes (such as [27]) is an interesting research issue.

## VI. EXPERIMENTS AND RESULTS

The experiments aim to study the impact on the routing algorithms by (1) the distribution patterns of the values of SSVs in the value space (*SSV Value Distribution*), (2) the behaviour or access pattern of ALPs on SSVs, (3) the physical distribution of SSVs in the CLP tree (*SSV Distribution Pattern*), and (4) the distribution of the values of SSVs in each individual CLP<sup>6</sup>. Factor 1 and 2 are at application level, related only to the agents and their environment, whereas factor 3 and 4 are at simulation level and can be parameterised. Moreover, certain additional factors have particular influence on the range-based routing approach, namely (5) the *ratio* of number of updates to number of range queries, (6) the *fluctuation* between the updated value and the original value, and (7) the *granularity of value segments*.

The environment for the experiments is set as a 2-dimensional space containing 6,200 objects of 8 types. SSVs are defined as the X-pos or Y-pos of any object in the environment (12,400 SSVs of 16 types), and they have numerous value distributions. On initialisation, the 64 agents

<sup>5</sup> The overall impact of combining routing and load management will be benchmarked on the PDES-MAS framework using realistic cases.

<sup>6</sup> This is different to the aforementioned value distribution parameter. For example, in a scenario the value ranges of SSVs in all CLPs are very close, and in another scenario the value ranges are significantly distinct from one CLP to another.

are distributed in a random, uniform manner over the space. The default settings of other parameters for the experiments are summarised in Tables I and II. The step size per dimension and the target distance conform to normal distributions, ( $\mu = 2.0$ ,  $\sigma = 1.0$ ) and ( $\mu = 5.0$  and  $\sigma = 2.0$ ), respectively. The diameter of an agent’s region of interest is set to 2.0. For example, an agent at the Pos-X = 50 could query a range of [49, 51) for all types associated with this dimension. The parameters roughly represent average values between worst and best cases for each routing approach. This allows us to observe the behaviour of each algorithm in a relatively small parameter space. The parameter space region that we investigated contains both kinds of setups: those in which the address-based routing is preferable, and those in which range-based routing is better (as illustrated for instance in Fig. 10, 11 and 21). In each of the following experiments, we always adjust one parameter while applying default values to the rest. This tends to explore the “general” scenarios in reality, and avoids biased analysis due to only using special cases with extreme settings. Each agent randomly generates eight requests for any type of SSV in each step. The simulation executes 300 time steps, which allows the simulated system to evolve significantly.

TABLE I  
SUMMARY OF THE DEFAULT APPLICATION LEVEL PARAMETERS

Name of Parameters	Value
Time steps	300
Events to be generated per agent per time step	8
Step size of all agents per dimension	$\mu = 2.0, \sigma = 1.0$
Target distance of all agents in steps	$\mu = 5.0, \sigma = 2.0$
Agent’s range of interest	2.0
Environment of the Agents	100 x 100 Torus

TABLE II  
SUMMARY OF THE DEFAULT SIMULATION LEVEL PARAMETERS

Name of Parameters	Value
Depth of the CLP Tree	4
Number of client ALPs to each server CLP	4 ALPs per server CLP
Root Imbalance	0
CLP Fluctuation	1
Number of segments for routing algorithms	100

The physical distribution of SSVs can be adjusted with two parameters, namely *root imbalance*

*ance* and *CLP fluctuation*. Root imbalance describes the percentage of additional SSVs hosted by the root CLP comparing to the rest. When root imbalance = 0, SSVs are distributed evenly on all CLPs. When root imbalance = 1, all SSVs are concentrated on the root CLP. Let  $n$  be the number of CLPs,  $r \in [0, 1]$  the root imbalance, and  $s_i \in [0, 1]$  the share of SSVs hosted by  $CLP_i$ ,

i.e.,  $s_2 = 0.6$  means that 60% of all SSVs are hosted by  $CLP_2$ . We can then define  $s_1$  (the share of

SSVs on the root CLP) as  $\frac{1+r(n-1)}{n}$  and  $s_i$  with  $i = 2, \dots, n$  as  $\frac{1-r}{n}$ , so that

$\sum_{i=1}^n s_i = \frac{1+r(n-1)}{n} + (n-1) \cdot \frac{1-r}{n} = 1$ . The CLP fluctuation constraints the maximum difference

between the greatest and smallest value of SSVs of the same type on an individual CLP. For example, suppose CLP fluctuation = 0.05 and the designated type of SSVs has a value space from 0 to 100, then the difference of these SSVs' value on the CLP is not greater than  $0.05 \times 100 = 5$ : A CLP may host two SSVs of type X-pos with values 80 and 82, but another SSV of type X-pos, with value 75, cannot be hosted by the same CLP, because  $82 - 75 > 5$ . Since SSVs do not migrate but have dynamic values, this condition holds only for the initial state. The experimental results are reported in terms of routing cost and accuracy.

The routing cost is measured using two metrics: the *number of messages* and the *number of hops to be traversed* in resolving each access to SSVs. The number of messages is the number of all messages that are generated by the routing algorithm in order to resolve a query. We regard the transmission of information from one CLP to another as a single message. If the same information is propagated along several CLPs, it is counted as multiple messages. Hence, the number of messages is a measurement of the overall bandwidth consumption. The number of hops is the maximum number of messages that had to be sent sequentially until the request could be resolved. This means, that the number of hops corresponds to the maximal path length from the

ALP generating the request to a CLP which had to be contacted, multiplied by 2 (for the query and the corresponding response). This is very similar to the notion of critical paths in parallel computing. Hence, the number of hops is a measurement of the overall latency. The two metrics for the routing algorithms are denoted by the variables  $range\text{-based}_{messages}$ ,  $range\text{-based}_{hops}$  and  $address\text{-based}_{messages}$ ,  $address\text{-based}_{hops}$  for the range-based and address-based algorithm respectively.

In order to calculate the accuracy of the routing algorithms, the minimal number of messages and hops ( $opt_{messages}$  and  $opt_{hops}$ ) are computed, which constitute the absolute limit for optimising the cost of any routing algorithm. For instance, Fig. 1 illustrates the querying of two SSVs, with the smallest set of CLPs and connections for this query highlighted. Each connection needs to transmit two messages (one request and one response), thus the total number of messages is 10. The maximum number of hops is recorded as 8 (for reading SSV<sub>1</sub> on the left); this is because messages have to be sent sequentially (in any path) until reaching the target and thus the latencies in other concurrent search paths are masked. The communication cost between the ALP and its server CLP is considered negligible. The ratio of the minimal number of messages (or hops) to the number of messages (or hops) is defined as the *accuracy* of routing algorithms. For example, the “message accuracy” of the range-based algorithm is defined as

$accuracy_{messages}^{range\text{-based}} = \frac{opt_{messages}}{range\text{-based}_{messages}}$ . Likewise, the “hop accuracy” for ranged-based algo-

rithm and the message and hop accuracy of the address-based algorithm are denoted by

$accuracy_{hops}^{range\text{-based}}$ ,  $accuracy_{messages}^{address\text{-based}}$  and  $accuracy_{hops}^{address\text{-based}}$  respectively.

#### A. Effect of Agent’s Environment (SSV Value Distribution)

Several non-uniform distributions of SSV values have been used while keeping default values

for all other parameters. Values have been assigned to SSVs in a round-robin fashion by one of three normal distributions (see Fig. 11(B)). The mean values of normal distributions are  $16\frac{2}{3}$ , 50 and  $83\frac{1}{3}$ . The deviation  $\sigma$  is varied from 0 to 10, so that the SSV value distribution changes from highly concentrated to highly scattered.

The effects of value distribution on the cost and accuracy of routing algorithms as compared to the optimal cost are reported in Fig. 12 and Fig. 13 respectively. Evidently, with values of SSVs distributed more sparsely, routing becomes more costly. In terms of number of hops, the range-based algorithm performs much better, by closely approaching  $opt_{hops}$ . When SSV values are scattered enough, the number of hops in both algorithms converges to  $opt_{hops}$ . In terms of number of messages, the performance of both algorithms deteriorates as the degree of dispersion of SSVs increases. The  $accuracy_{hops}$  of both algorithms converge to 1 while the maximum  $accuracy_{messages}$  for both algorithms only approximates 0.8. In all situations, the range-based algorithm is likely to incur less overhead for routing range queries.

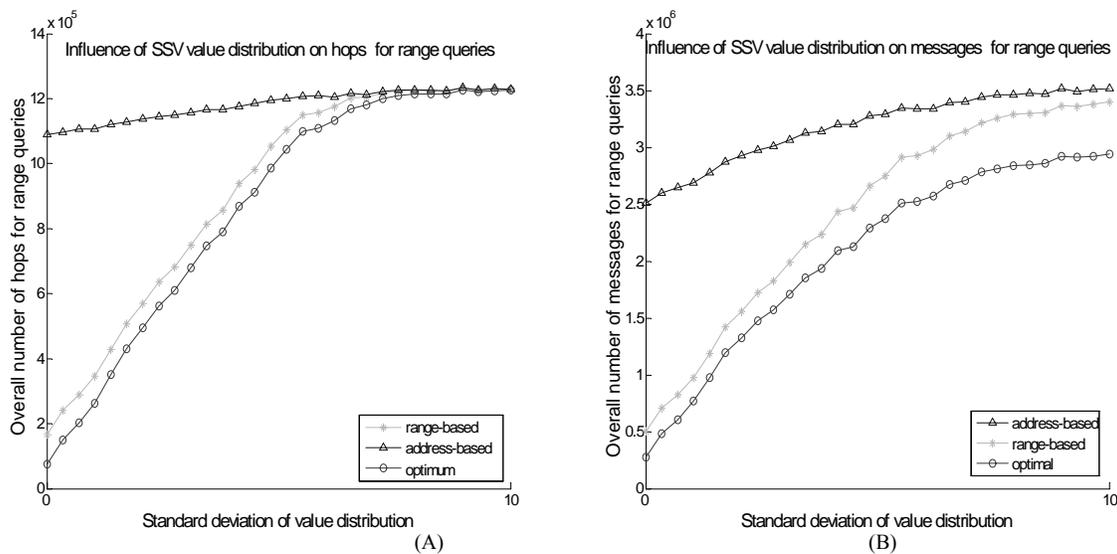


Fig. 12. Influence of Overall SSV Value Distribution on Routing Cost

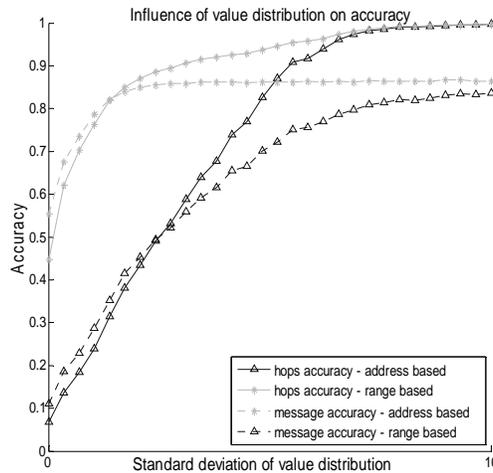


Fig. 13. Influence of Overall SSV Value Distribution on Accuracy of Routing

Similar experiments have also been performed to measure the accuracy of making ID queries.

The address-based algorithm always achieves optimal results (accuracy = 1) while the accuracy of the range-based algorithm is quite low, about 0.36. This is because the range-based algorithm does not store the location of an individual SSV and requires exhaustive range searches on SSV IDs.

### B. Effect of SSV Value Distribution Pattern

A set of experiments have been performed to examine how the distribution pattern of SSVs in the simulation infrastructure (the CLP tree) affects the performance of routing algorithms. This subsection reports the effect of two simulation level factors: (a) the physical distribution of SSVs on the CLP tree, and (b) the distribution of SSVs' values on each individual CLP, given the behaviour of agents remains the same. The two relevant parameters, *root imbalance* and *CLP fluctuation*, are varied between 0..1 and 0.05..1 respectively.

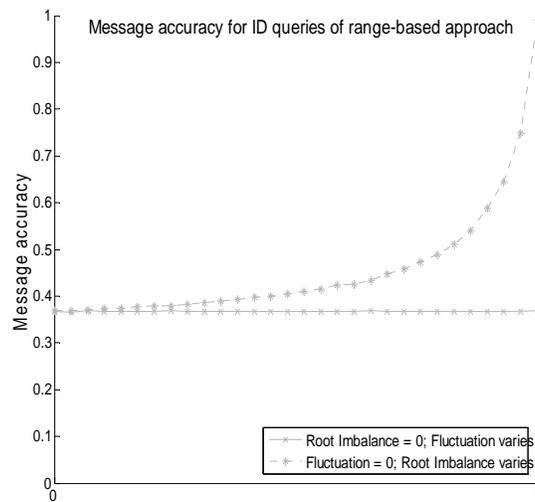


Fig. 14. Message Accuracy for ID Query

Fig. 14 gives the accuracy of the range-based algorithm for ID queries in terms of messages. The SSV value distribution in a CLP does not affect the routing of ID queries at all. However, concentration of SSVs on fewer CLPs will dramatically improve the accuracy of routing ID queries using range-based algorithm.

Fig. 15 illustrates the impact of the SSVs' physical distribution pattern on the cost of routing range queries using the two routing algorithms. When SSVs are distributed on all CLPs uniformly (root imbalance = 0),  $range\text{-based}_{hops}$  and  $address\text{-based}_{hops}$  are nearly equal to  $opt_{hops}$  while  $range\text{-based}_{messages}$  and  $address\text{-based}_{messages}$  are very close but still much greater than  $opt_{messages}$ . With the increase of root imbalance, the routing cost gradually decreases. The range-based algorithm adapts much better than the address-based algorithm. When root imbalance = 1, all SSVs locate at the root CLP and it makes no difference to either algorithm; this extreme case reflects the use of a single centralised CLP.

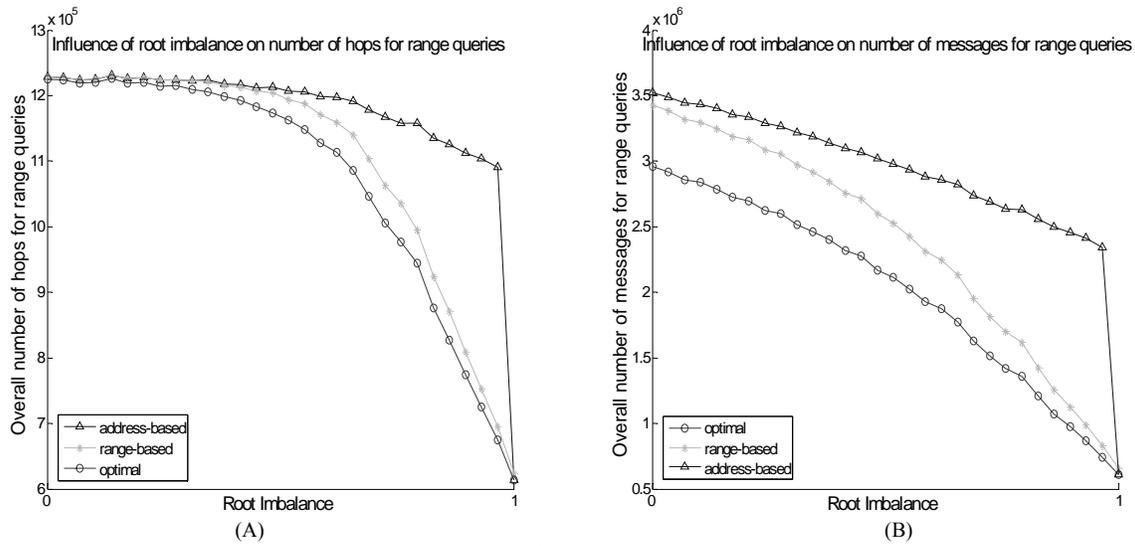


Fig. 15. Influence of the SSV Distribution Pattern on the Routing Cost

Fig. 16 presents the cost of performing range queries against the value distribution on each CLP. The results are similar to those obtained in Fig. 12, namely the value distribution of all SSVs and the value distribution of SSVs on each individual CLP both have significant influence on the routing cost involved in range queries.

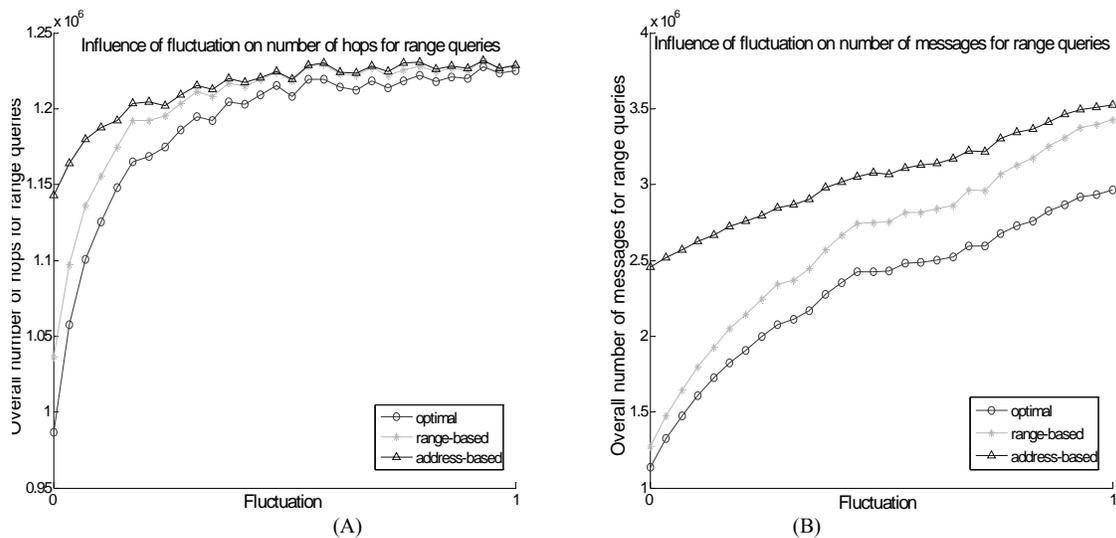


Fig. 16. Influence of SSV Value Distribution on Individual CLPs on Routing Cost

The accuracy of the routing algorithms against the SSVs' physical distribution pattern is illustrated in Fig. 17. These results further indicate that the range-based algorithm adapts better to the SSV's distribution pattern. The latencies (number of hops) incurred by both algorithms are simi-

lar; however the range-based algorithm generates much less communication traffic (number of messages). Furthermore, comparing to Fig. 13, the value distribution on an individual CLP has a less significant impact than the overall value distribution.

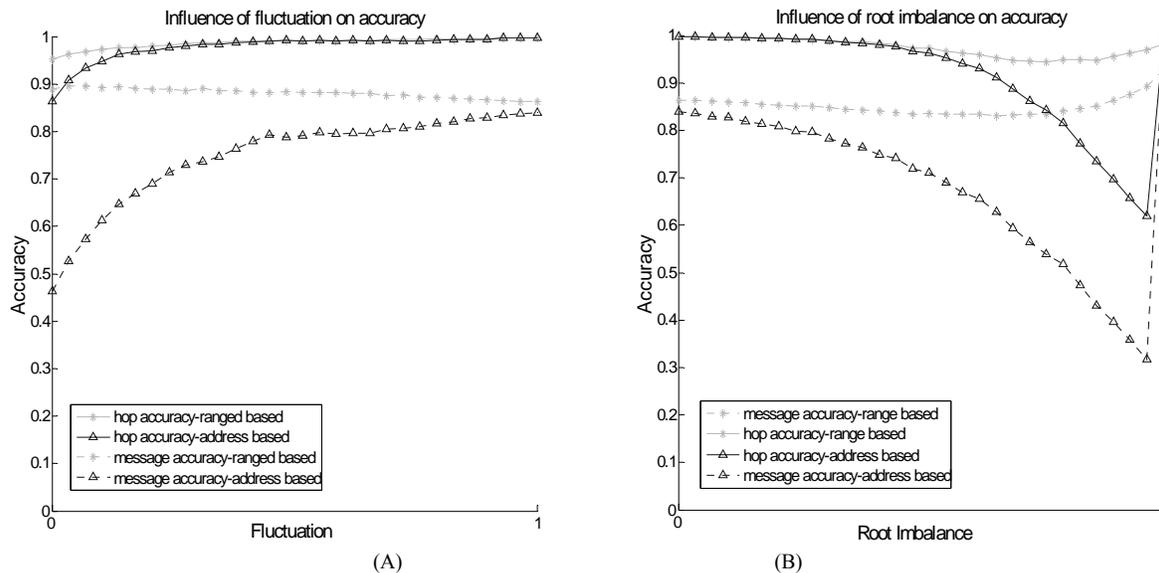


Fig. 17. Influence of SSV's Physical Distribution Pattern on the Accuracy of Routing Algorithms

### C. Further Analysis of the Range-based Approach

The complexity of the range-based approach calls for further investigation. Three additional key parameters may affect the performance of range-based routing: (a) the *ratio* of the number of updates to number of range queries, (b) the *fluctuation* between the updated value to the original value, and (c) the *granularity of segments*.

#### 1) Effect of SSV Update (Application Level)

An agent may randomly update any SSV whose value is in its region of interest. In the meta-simulation, the SSV's value is updated using an absolute offset (to the old value), which is randomly set conforming to uniform distribution with lower bound 0 and upper bound varied between 0.01 and 5. Furthermore, the probability that an agent generates an update query is set between 0.01 and 0.5.

Fig. 18 gives an overall picture of the correlation between offsets of updated SSV values and

the ratio of update queries to the messages (updates) needed for accomplishing update queries. Three specific ratios of update queries are highlighted. The results indicate that the increase of both the offset and the ratio of update queries lead to the generation of more update messages. Compared to the number of overall messages obtained in the previous experiments (see Fig. 12, 15, 16), the number of update messages is very small. Although this number depends on the specific configuration of the experiments, it still suggests that the overhead incurred by update queries tends to be negligible in an environment with heavy load and traffic.

Fig. 19 illustrates the effect of the ratio of update queries to the routing cost and the breakdown of routing cost involved in update queries and range queries. This ratio has a significant impact to the number of messages and hops required by both query types. Additionally, Fig. 19(A) shows that basically the address-based algorithm outperforms the range-based one in situations with frequent update queries. Although the address-based approach generates more messages for range queries, it still incurs fewer messages in total due to the benefit in update queries. Fig. 19(B) indicates that the number of hops is linear to the ratio of update queries and the candidate routing algorithms do not make considerable difference in this aspect.

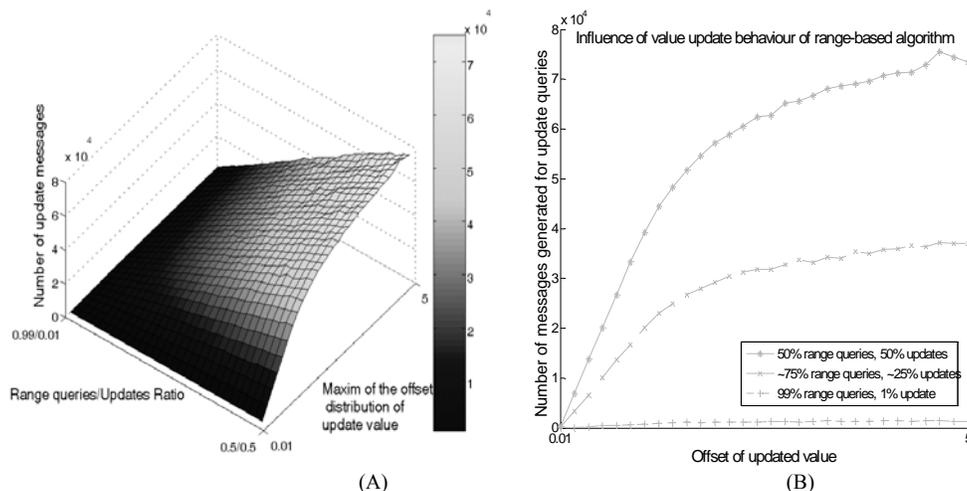


Fig. 18. Agent Behaviour's Influence on Number of Messages for Update Query

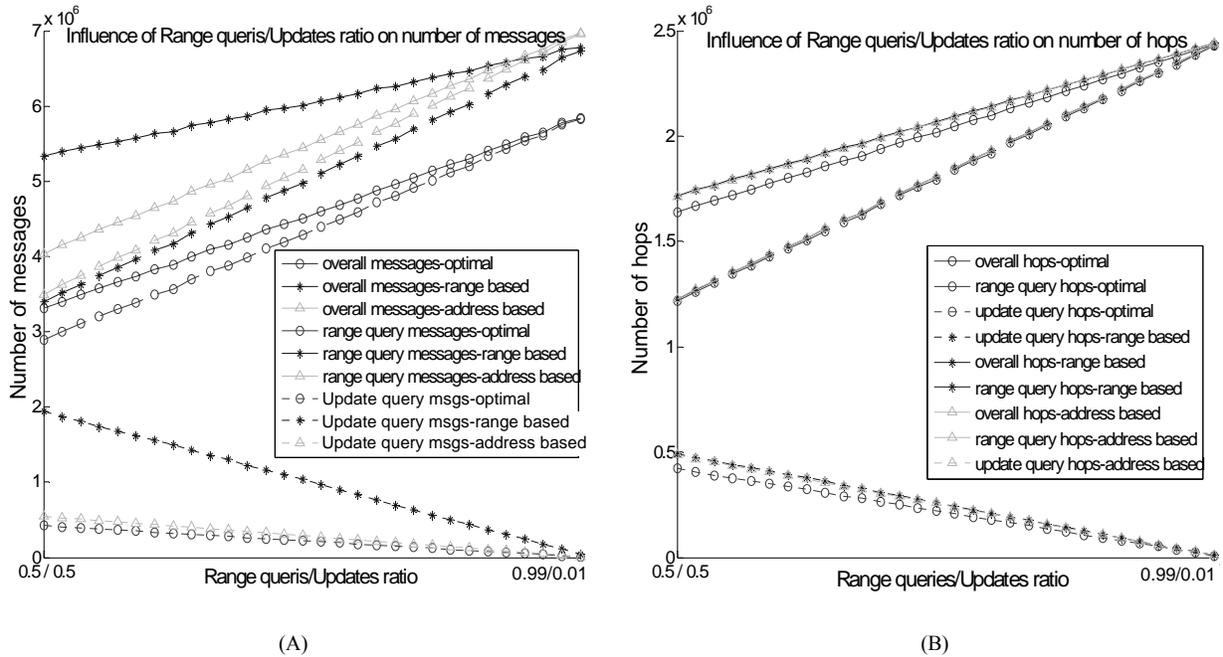


Fig. 19. Influence on the Ratio of Update Queries to Routing Cost

## 2) Effect of Granularity of Segments (Simulation Level)

A set of experiments have investigated the role of the granularity of attribute value range segmentation for the range-based algorithm. The number of segments used to describe an attribute value range varies between 1 and 200, with a larger number implying a more precise attribute value range description.

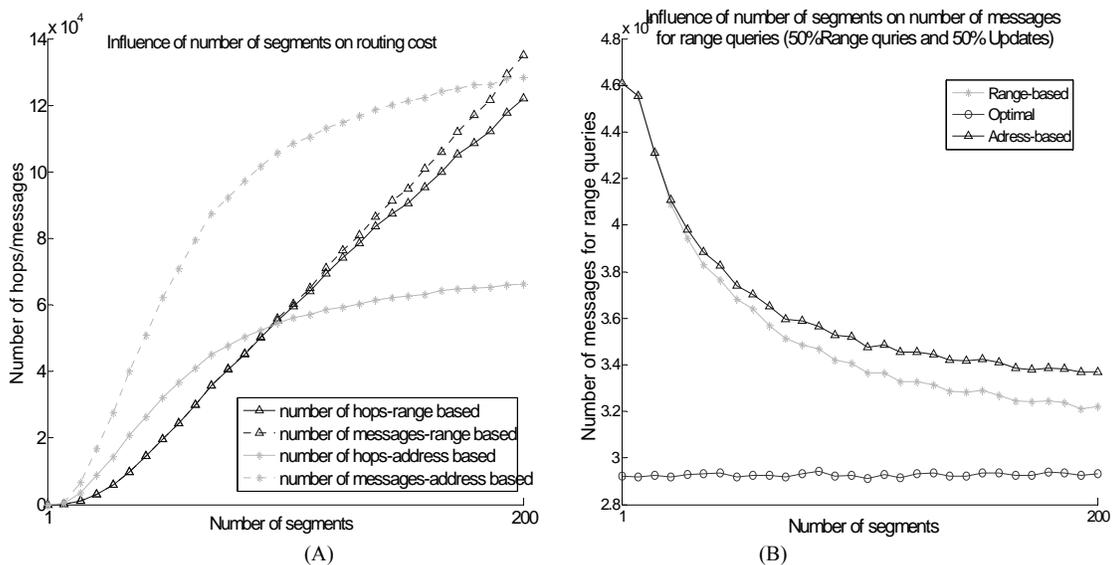


Fig. 20. Influence of Segment's Granularity on the Routing Cost

The results are reported in Fig. 20. The total (for both range queries and updates) number of messages required by the address-based algorithm is larger in most cases. As the number of segments increases the address-based algorithm latency stabilises, as queries have to be sent only to neighbours. When the number of segments becomes large enough, the range-based algorithm tends to update CLPs in a rather wide extent, an action which deteriorates its performance. Fig. 20(B) illustrates the relationship of routing range queries to the granularity of segmentation. As the segment size decreases, the number of messages required to satisfy the range query is dramatically reduced, because the accurate description of the value ranges results in a more precise routing.

Fig. 21 presents the overall difference (range-based - address-based) of the two algorithms in terms of (total) number of messages. Negative values denote the space for achieving better performance for adopting the range-based approach, and vice versa for adopting the address-based one.

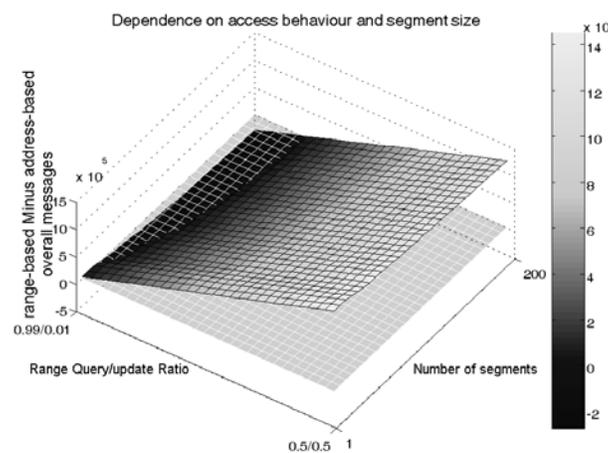
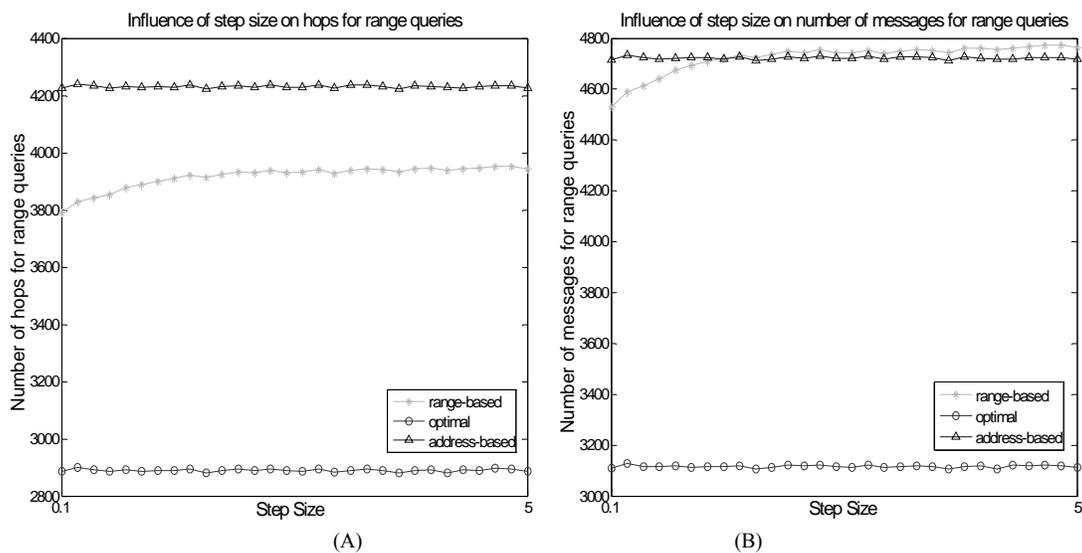


Fig. 21. Difference in Message Numbers using Different Routing Algorithms

### 3) Effect of Agent's Behaviour

In the experiments presented in the previous sections, the environment is heavily populated (64 agents randomly situated in a  $100 \times 100$  surface and 12,400 SSVs distributed either sparsely or

concentrated in 9 central points as shown in Fig. 11), a feature that renders the investigation of the effect of agents' behaviour a difficult task as, irrespective of how far agents move (regionally or globally), there are always plenty of SSVs for them to access in their region of interest. For this investigation, a thinly populated environment with 4 agents and 1,550 SSVs<sup>7</sup> of two types has been used; the size of CLP tree has been reduced accordingly (depth = 2). In this set of experiments, an agent's behaviour pattern can be characterised by its activity scope and the speed at which it moves. For each agent, its range of interest is 5.0 and its stepsize is varied between 0.1 (agent moves regionally) to 5.0 (agent moves globally). Two different target distances, (1, 0) and (5, 2), are tested to mimic different movement speeds. The rest of the parameters have been set to the default values.



<sup>7</sup> This number represents 1/16th of the number of SSVs (12,400) in the original experiments, to reflect the number of agents used (4=64/16)

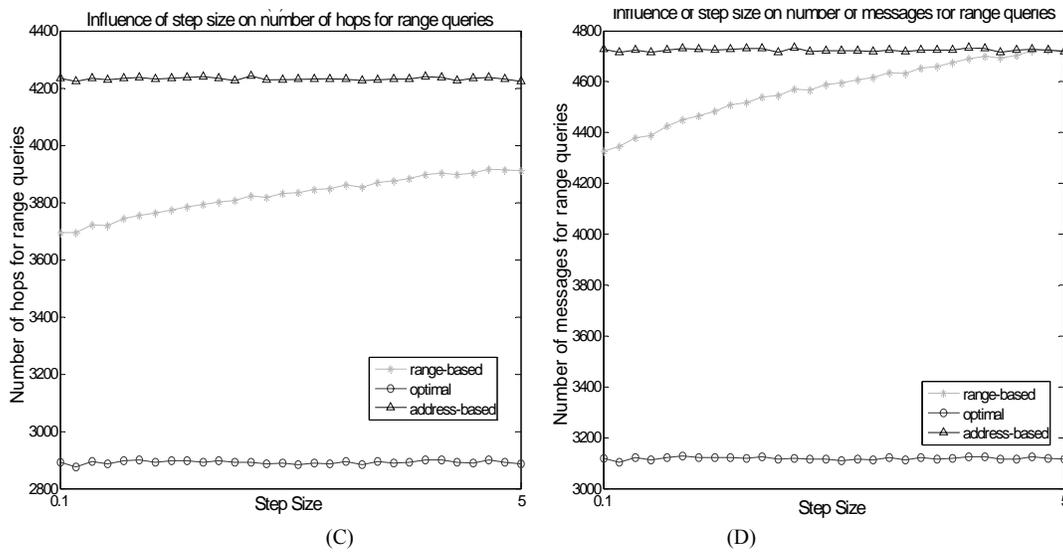


Fig. 22. Effect of Agent's Access Pattern

Fig. 22 reports the effect of the size of an agent's activity territory in terms of number of hops and messages. Fig. 22(A)&(B) report results for "fast" agents and (C)&(D) for "slow" agents. Evidently, the address-based approach incurs more latency while its overall cost is almost neutral to the agent's moving pattern. When an agent moves more globally and faster, the range-based algorithm becomes more costly as the agent generates queries that differ very much from each other, so that many updates occur.

## VII. RELATED WORK

The distributed simulation of agent-based systems has been the focus of several recent papers including [1][20][23][24][39][13][33][34], but none of this work has considered the efficient distribution and access of the simulation's shared state space.

Routing has been heavily studied in the area of computer networks [38], however the issue of range queries does not arise in this context. In [36] Steen et al. proposed a scalable location service mechanism for locating mobile objects in distributed environments. Their approach is address-based, which binds an object's name to one or more addresses where the object can be contacted. The location service is designed to handle objects with arbitrary migration patterns.

Using a hierarchical search tree, where each new object is registered, the leaf nodes store the addresses. By registering the contact address in the smallest region in which the object is moving, the approach only requires searching an extremely short path to locate randomly migrating objects.

The issue of range queries has received considerable attention in the area of distributed spatial databases, e.g., [2][9][12][26][28][35], although the work in this area tends to focus on the optimisation of matching algorithms. In the area of context-aware computing, location dependent information services (LDIS) address the problem of mobile query sources, however, typically they do not have to cope with highly dynamic data distributions, because the data can be distributed according to its geographical validity in a static fashion [17].

The problem of coordinated access to shared data has also been studied in the context of cache consistency protocols for distributed and multiprocessor systems (e.g., [22][25][37]), but such systems are based on data replication and therefore the solutions devised are not directly applicable to the PDES-MAS architecture.

In the Modelling and Simulation community, the majority of approaches focus on the use of multicast rather than unicast peer-to-peer (P2P) overlays as described here. Notable contributions by Rak et al [31], Boukerche et al [6], and Berrached et al [5] all used variations on the theme of mapping a set of multicast groups on to the cells of an n-dimensional grid. Range queries here are mapped to some subset of the cells (groups) to which the querying node subscribes. Alternative mappings such as that proposed by Morse [15] reduce redundant traffic, but increase the computational complexity.

The problem of how to optimally distribute and retrieve both data (publications) and queries over the data (subscriptions) in a P2P overlay has been studied extensively in the field of Content

Addressable Networks (CANs) such as the Distributed Hash Table (DHT) [32]. Approaches concentrating on serving contiguous range queries in several dimensions with minimum latency have been presented by both Barambe et al [3][4] and Ganesan et al [11]. These systems are designed for real time applications, such as P2P network multiplayer games, rather than asynchronous logical time systems like parallel simulations.

### VIII. CONCLUSIONS AND FURTHER WORK

This paper has identified the efficient data accessing as a key issue to optimising the execution of distributed simulations of complex agent-based systems. It has presented two different approaches to address this problem, namely an address-based approach, which locates data according to their address information, and a range-based approach, whose operation is based on looking up attribute value range information along the paths to the destinations.

To facilitate the analysis of the proposed approaches, the paper adopted a meta-simulation framework to analyse query performance on distributed and dynamic data whose location and properties are changing constantly and unpredictably. This is an open challenging problem and particularly difficult to tackle using pure deterministic or analytic approaches. Although the proposed algorithms were analysed in the context of the PDES-MAS framework, the results obtained are valid for any system with similar characteristics, such as distributed systems based on peer servers. The main conclusions that can be drawn from the experimental results obtained can be summarised as follows:

- The data value distribution has a vast impact on the performance of both approaches. When the data values are distributed sparsely enough, both algorithms can minimize the latency of routing to the optimal case, whereas they tend to generate a large number of messages otherwise (~20% more than the optimal case).

- The data distribution pattern is also a decisive factor to the performance of routing. The more evenly data is distributed, the closer both approaches approximate the optimal case. The smaller the fluctuation of data values maintained in a particular node, the more accurate the routing.
- Granularity of segmentation can considerably affect the routing of update queries. The address-based approach is superior to the range-based one when segments are very precise. In range-based routing, the routing accuracy is adapted. However, a precise segmentation leads to large overhead in dealing with update queries.

The meta-simulation analysis approach described in this paper has proved to be a powerful tool to analyse the impact of the model abstract characteristics to the efficiency of the proposed algorithms. However, as in every simulation exercise, the problem of verification and validation of the simulation and the reliability of the results remains. Most importantly, as any meta-simulation approach, it provides only an idealistic performance projection of the algorithms and does not take into account the specific characteristics of any particular implementation or agent-based model. A more realistic performance evaluation of the proposed routing algorithms and an analysis of their scalability and their computational complexity requires their implementation and integration in the PDES-MAS kernel and experimentation with more realistic benchmarks. Work in this direction has already commenced.

Another challenging issue to be addressed in the future is the relationship between the proposed routing algorithms and the load management (data migration) and synchronisation mechanisms of the PDES-MAS kernel (described in [27] and [18][19] respectively).

#### ACKNOWLEDGEMENTS

This work was supported by the EPSRC under Grant No. GR/R45338/0.

## REFERENCES

- [1] John Anderson. A generic distributed simulation system for intelligent agent design and evaluation. In Proceedings of AI, Simulation and Planning In High Autonomy, Systems, 2000.
- [2] N. An, J. Jin, and A. Sivasubramaniam. "Toward an Accurate Analysis of Range Queries on Spatial Data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 15, issue 2, pp. 305-323. 2003.
- [3] A. Barambe, M. Agrawal, and S. Shesan. "Mercury: Supporting Scalable Multi-Attribute Range Queries," *Proceedings of the ACM SIGCOMM*, 2004.
- [4] A. Barambe, J. Pang and S. Shesan. "A Distributed Architecture for Multiplayer Games," Technical Report, CMU-CS-05-112. Carnegie-Mellon University, 2006.
- [5] A. Berrached, M. Beheshti, O. Sirisaengtaksin and A. deKorvin. "Approaches to Multicast Group Allocation in HLA Data Distribution Management," *Proceedings of the 1998 Spring Simulation Interoperability Workshop*, 1998.
- [6] A. Boukerche, A. J. Roy, and N. Thomas. "Dynamic Grid-Based Multicast Group Assignment in Data Distribution Management," *Proceedings of the Fourth IEEE International Workshop on Distributed Simulation and Real-Time Applications*, 2000.
- [7] J. M. Epstein and R. L. Axtell. *Growing Artificial Societies: Social Science From the Bottom Up*, Brookings Institution Press, 1996.
- [8] R. Ewald, "Simulation of Load Balancing Algorithms for Discrete Event Simulations", Master Dissertation, School of Computer Science, University of Rostock, Germany. [in preparation]
- [9] C. Faloutsos and I. Kamel. "Beyond Uniformity and Independence: Analysis of R-Trees Using the Concept of Fractal Dimension," *Proceedings of the thirteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pp. 4-13. June 1994.
- [10] Richard M. Fujimoto, *Parallel and Distributed Simulation Systems*, Wiley Series on Parallel and Distributed Computing, 2000.
- [11] P. Ganesan, B. Yang, and H. Garcia-Molina. "One Torus to Rule them All: Multi-Dimensional Queries in Peer-to-Peer Systems," *Proceedings of the 7th International Workshop on the Web and Databases*. 2004.
- [12] Q. He, M. H. Ammar, G. Riley, H. Raj, and R. Fujimoto. "Mapping Peer Behavior to Packet Level Details: A Framework for Packet-level Simulation of Peer-to-Peer Systems," *Proceedings of the 11th International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, 2003Y. Ioannidis, "Query Optimization", *ACM Computing Surveys*, symposium issue on the 50th Anniversary of ACM, vol. 28, no. 1, pp. 121-123. March 1996.
- [13] Jan Himmelspach, Roland Ewald, Stefan Leye, and Adelinde M. Uhrmacher. Parallel and distributed simulation of parallel DEVS models. In Proceedings of the SpringSim '07, DEVS Integrative M&S Symposium, pages 249–256. SCS, 2007.
- [14] N. R. Jennings and M. Wooldridge, "Applications of intelligent agents," in *Agent Technology: Foundations, Applications, Markets*, N. R. Jennings and M. Wooldridge, Eds., pp. 3–28. Springer-Verlag, 1998
- [15] K. Morse, L. Bic, M. Dillencourt, and K. Tsai. "Multicast Grouping for Dynamic Data Distribution Management," *Proceedings of the 31st Society for Computer Simulation Conference*. 1999.
- [16] J. Liu, D. Nicol, B. Premore, and A. Poplawski. "Performance Prediction of a Parallel Simulator," *Proceedings of the 13th Workshop on Parallel and Distributed Simulation (PADS'99)*, pp. 156-164. 1999.
- [17] D.L. Lee, J. Xu, B. Zheng, and W. Lee. "Data Management in Location-Dependent Information Services," *IEEE Pervasive Computing*, vol. 1, issue 3, pp. 65-72. 2002.
- [18] M. Lees, B. Logan, and G. Theodoropoulos. "Adaptive optimistic synchronisation for multi-agent simulation," In D. Al-Dabass, editor, *Proceedings of the 17th European Simulation Multiconference*, pp. 77–82. 2003.
- [19] M. Lees, B. Logan, D. Chen, T. Oguara, and G. K. Theodoropoulos. "Decision-Theoretic Throttling for Optimistic Simulations of Multi-agent Systems," *Proceedings of the Ninth IEEE International Workshop on Distributed Simulation and Real-Time Applications*, pp. 179-186. October 2005.
- [20] M Lees, B Logan, GK Theodoropoulos. 2007. Distributed Simulation of Agent-based Systems in HLA, *ACM Transactions on Modelling and Computer Simulation*, Vol. 17, 3, Available at: <http://doi.acm.org/10.1145/1243991.1243992> ISSN: 1049-3301.
- [21] B. Logan and G. K. Theodoropoulos. "The distributed simulation of multi-agent systems," *Proceedings of the IEEE*, vol. 89, no. 2, pp. 174-186. 2001.
- [22] C. Morin and I Puaut, "A Survey of Recoverable Distributed Shared Virtual Memory Systems", *IEEE Transactions on Parallel and Distributed Systems*, September 1997, Vol. 8, No. 9 pp. 959-969
- [23] Robert Minson and Georgios Theodoropoulos. Distributing RePast agent based simulations with HLA. In Proceedings of the 2004 European Simulation Interoperability Workshop, Edinburgh, 2004.
- [24] R. Minson, G. Theodoropoulos, Distributing RePast agent-based simulations with HLA, *Computation and Concurrency Practice and 73*, Experience Journal, Wiley, 26 pages [in press].
- [25] S.V. Nagaraj, *Web Caching and Its Applications*, Springer, 2004, 262 pages, ISBN: 1-4020-8049-2
- [26] S. Ndiaye, M. Tsangou, M. Seck, and W. Litwin. "Range Queries to Scalable Distributed Data Structure RP\*," *Proceedings of the Fifth Workshop on Distributed Data and Structures*, 2003.
- [27] T. Oguara, D. Chen, G. K. Theodoropoulos, B. Logan, and M. Lees. "An Adaptive Load Management Mechanism for Distributed Simulation of Multi-agent Systems," *Proceedings of the Ninth IEEE International Workshop on Distributed Simulation and Real-Time Applications*, pp. 179–186. October 2005.
- [28] B. Pagel, H. W. Six, H. Toben, and P. Widmayer. "Towards an Analysis of Range Query Performance in Spatial Data Structures," *Proceedings of the twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pp. 214-221. May 1993.
- [29] K. S. Perumalla, R. M. Fujimoto, P. J. Thakare, S.Pande, H. Karimabadi, Y. Omelchenko, and J. Driscoll. "Performance Prediction of Large-Scale Parallel Discrete Event Models of Physical Systems," *Proceedings of Winter Simulation Conference 2005*, 2005.
- [30] B. Rajive, E. Deelman, and T. Phan. "Parallel Simulation of Large-Scale Parallel Applications," *The International Journal of High Performance Computing Applications*, vol. 15, no.1, pp. 3-12. 2001.
- [31] S. J. Rak, M. Salisbury, and R. S. MacDonald, "HLA/RTI Data Distribution Management in the Synthetic Theater of War," *Proceedings of the Fall 1997 DIS Workshop on Simulation*, 1997.
- [32] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. "A Scalable Content-Addressable Network," *Proceedings of the ACM SIGCOMM*, 2001.

- [33] Patrick Riley. MPADES: Middleware for parallel agent discrete event simulation. In Gal A. Kaminka, Pedro U. Lima, and Raul Rojas, editors, *RoboCup-2002: The Fifth RoboCup Competitions and Conferences*, number 2752 in *Lecture Notes in Artificial Intelligence*. Springer Verlag, Berlin, 2003.
- [34] Patrick Riley and George Riley. SPADES — a distributed agent simulation environment with software-in-the-loop execution. In S. Chick, P. J. Sanchez, D. Ferrin, and D. J. Morrice, editors, *Winter Simulation Conference Proceedings*, 2003.
- [35] Hanan Samet, *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, 1990
- [36] M. V. Steen, F. Hauck, P. Homburg and A. S. Tanenbaum, "Locating Objects in Wide-area Systems," *IEEE Communications Magazine* vol. 36, no. 1, pp. 104–109. 1998.
- [37] Per Stenström, "A Survey of Cache Coherence Schemes for Multiprocessors", *ACM Computer* , Volume 23 , Issue 6 (June 1990), Pages: 12 - 24 , ISSN:0018-9162
- [38] A. S. Tanenbaum, *Computer Networks*, Fourth Edition. ISBN 0-13-066102-3, Prentice Hall, USA, 2003.
- [39] F. Wang, S.J. Turner, and L. Wang. Integrating agents into HLA-based distributed virtual environments. In *4th Workshop on Agent-Based Simulation*, pages 9–14, Montpellier, France, 2003.
- [40] D. Wyens, H. Parunak, F. Michel, T. Holvoet and J. Ferber. "Environments for Multi-agent Systems, State-of-art and Research Challenges," *Lecture Notes in Computer Science*, vol. 3374, 1-48. 2005.
- [41] Roland Ewald, Dan Chen<sup>1</sup>, Georgios Theodoropoulos, Michael Lees., Brian Logan, Ton Oguara, Adelinde M. Uhrmacher, "Performance Analysis on Shared Data Accessing Algorithms for Distributed Simulation of Multi Agent Systems", presented at the 20th *ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS 2006)*, May 23–26, 2006, Raffles Hotel, Singapore.