

# THE ROLE OF TRUST IN DISTRIBUTED DESIGN

NIEK J. E. WIJNGAARDS, HIDDE M. BOONSTRA AND FRANCES M. T. BRAZIER

Intelligent Interactive Distributed Systems Group,  
Faculty of Sciences, Vrije Universiteit Amsterdam,  
de Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands  
Phone: +31 - (0)20 - 444 7634. Fax: +31 (0)20 - 444 7653.  
Email: {NJE.Wijngaards, HM.Boonstra, FMT.Brazier}@few.vu.nl

URL: <http://www.iids.org/>

**Abstract.** Automated support of design teams, consisting of both human and automated systems, requires an understanding of the role of trust in distributed design processes. By explicating trust, an individual designer's decisions become better understood and may be better supported. Each individual designer has his or her private goals in a co-operative design setting, in which requirement conflicts and resource competitions abound. There are, however, also group goals that also need to be reached. This paper presents (1) an overview of research related to trust in the context of agents and design, (2) a computational knowledge-level model of trust based on the seven beliefs distinguished by Castelfranchi and Falcone (2000); and (3) an example of the use of the trust model in a specific design process, namely website design from the perspective of a single designer. The results are discussed in the context of distributed design in open systems.

**Keywords:** trust model, distributed design, knowledge-level, agents.

## 1. INTRODUCTION

In distributed design, individual designers work together to solve a specific design problem. Such group design processes differ from individual design processes (Dwarakanath & Blessing, 1996) in a number of ways: explication of design steps is often needed to facilitate interaction between the different parties involved, resulting in a larger number of alternatives being explored (see Cross et al. (1996)). Partial or complete automated support for such distributed design processes necessitates an understanding of such distributed processes, to be able to build models on which new systems can be based.

Each individual designer reasons explicitly about the situation in which the design process is performed: about his/her interpretation of a specific situation (Gero, 1998; Maher et al., 2000; Gero & Kannengiesser, 2002). Each designer has his or her own view of the world and other agents, and their environments, including assessments of their expertise, reliability, experience, trust, et cetera. Such assessments require reflection (Schön, 1983). A model of the types of reflection of individual designers in distributed design is presented in Brazier et al. (2001a).

This paper focuses on the role of trust in distributed design. Distributed design involves a number of participants, each with their own characteristics: e.g. expertise, experience, goals, and attitudes (e.g., Busby, 2001). Information acquired from different participants may be valued differently in terms of accuracy and trustworthiness, depending on the context in which it is acquired. Human participants in a distributed design setting often know whom they trust, and whose abilities they value, and when. This knowledge is not often made explicit while it does influence distributed design processes (i.e., the way in which members of a design team assess and incorporate each others' designs, objectives, and evaluations). These trust relations need to be made explicit to be able to acquire the models with which complete or partial design support systems can be developed.

Agents are a useful metaphor for designers. Agents, in this context, are defined as social, reactive, proactive and autonomous entities (Wooldridge & Jennings, 1995). This is supported by work done by, e.g., Lawson (1997) and Cross et al. (1996) in design: autonomy, co-operation and competition are the basic characteristics of designer agents to which they refer:

- A designer agent needs autonomy to exert control over its own processes, deliberate about co-operative work, its attitude, etc.
- A designer agent needs to be co-operative to jointly work with other designer agents, employing shared ontologies, shared protocols, shared communication languages, shared agreements, a shared model of design, etc.
- A designer agent needs to be competitive to deliberate about its own goals in relation to the goals of other agents and the aims of the design project, make (somewhat) selfish decisions to prevent overload, persuade other designers, etc.

Co-operation between agents is essential for the single function agents (SiFAs) of Grecu and Brown (1996). These agents, however, have very limited knowledge. Relatively few examples of the use of agents in design processes exist. In A-design (Campbell et al., 1998) agents are used to model evolutionary computation: designer agents are specialised in either creating, assessing, or removing solutions. As such, their role in a distributed process is not that of a human designer. This is the case for the agents used by McAlinden et al. (1998) to support information and knowledge handling in a design project. Collaboration between these agents is, however, minimal.

Co-operation in distributed or concurrent human design processes is subject to automation by facilitation (Boujut & Laureillard, 2002), e.g. by introducing tools to monitor progress and understanding (Hill et al., 2002), and to analyse participation (Simoff & Maher, 2000). Collaboration is often made explicit in agent-based design systems from an engineering perspective (e.g., Wilson & Shi, 1996; Lees et al., 2001; Anumba et al., 2002; Liu & Frazer, 2002; Zha, 2002), focussing on task-coordination without explicitly incorporating trust. An overview of design processes involving collaborative agents is provided by, for example, (Wang et al., 2002).

This paper first discusses the current state of art with respect to research on trust in Section 2. A computational model of trust, based on Castelfranchi & Falcone (2000) is introduced. This model is used in a knowledge-level analysis of trust relations in distributed website design. Two specific design projects<sup>1</sup> are analysed, in particular with respect to task delegation from the perspective of a single designer. Formalisation of the knowledge involved (including trust) provided a means to computationally simulate the processes involved, making it possible to evaluate the results. Section 3 presents this work. Section 4 discusses the results and indicates areas in which further research is required.

## 2. TRUST

This section presents a brief overview of research on trust in agent systems and design. Section 2.1 explores the notion of trust. Section 2.2 discusses research on trust related to agents and design. Section 2.3 discusses trust models. Section 2.4 introduces a new computational trust model based on Castelfranchi and Falcone (2000)'s trust model. This model is used to analyse the role trust plays in a distributed website design project described in Section 3.


---

<sup>1</sup> The authors thank Us Media ([www.usmedia.nl](http://www.usmedia.nl)) for their permission to use the data on these projects.

## 2.1 DEFINING TRUST

Trust is a complex, subjective concept with many definitions from various fields of research such as psychology, management & communications, sociology, economics and political sciences (McKnight et al., 1998; McKnight & Chervany, 2001). McKnight and Chervany (2001) and McKnight et al. (2002) specify a conceptual typology of high level trust and distrust concepts on the basis of literature research, combining trust definitions from various research areas. Five conceptual trust types are distinguished and can be used to guide researchers in choosing their definition of trust. The five types defined by McKnight et al. (2002) and the relations between these types are depicted in Figure 1.

- Disposition to Trust: general extent to which trust is placed in others.
- Institution-based Trust: truster beliefs favourable conditions are in place.
- Trusting Beliefs: extent to which truster beliefs characteristics of the trustee.
- Trusting Intentions: the truster is *willing* (intending) to depend on the trustee.
- Trust-related Behaviour: the truster *depends* on the trustee.



*Figure 1: Trust types and their relations (McKnight and Chervany, 2001).*

Trust is clearly about consequences related to risk and actions (e.g., Luhmann, 2000). McKnight et al. (1998) distinguish more precisely three essential elements in trust: (a) potential negative consequences, (b) dependency and (c) feelings of security. In this paper, trust is viewed as a combination of: trusting beliefs, trusting intentions, and trust-related behaviour. This is in accordance with Gambetta (2000): the subjective probability that one or more trustees will perform a particular action.

## 2.2 TRUST AND AGENTS

Whenever agents (human or automated) co-operate, compete, perform transactions, or engage in other interactions, trust plays a role (Falcone et al., 2001; Deutch, 1962; Castelfranchi & Falcone, 1998; Cambetta, 2000). In essence, all co-operative tasks include trust (Gambetta, 2000), as an individual reduces autonomy by increasing dependence on other individuals. Although not often recognized explicitly (e.g., Schön, 1983; Coates et al., 2000; Chao et al., 2002; Brazier et al., 2001b; Valkenburg & Dorst, 1998; Wang et al., 2002), this is also the case in distributed design (Brazier & Wijngaards, 2002). In distributed design, trust plays a role in delegation (e.g., Milewski & Lewis, 1997) and trust plays a role in assessing information sources (Hertzum & Pejtersen, 2000; Hertzum, 2002). Trust in other agents, trust in information, and trust in organisations, all play a role.

This paper focuses on an individual designer's trust in other agents. An individual designer in a distributed setting is a reflective designer (e.g., Schön, 1983; Atman et al., 1999; Adams et al., 2003); whose professional skills also involve interpersonal and project management skills (Lewis & Bonollo, 2002) sometimes assuming different roles (e.g., Sonnenwald, 1996). Different aspects of trust identified in the literature (e.g., for an overview, see Wang et al., 2002) include reputation (e.g., Craig & Zimring, 2000; Lang et al., 2002), knowledge about collaborators (e.g., Denton, 1996; Busby, 2001), delegation and co-ordination (e.g., Wilson & Shi, 1996; Tambe, 1997; Valkenburg & Dorst, 1998; Jensen et al., 2000; Pynadath & Tambe,

2002), risk (e.g., Valkenburg & Dorst, 1998), expectations (Gero & Kannengiesser, 2002), and judgements (Holt, 1997).

The trust an agent has in other agents is in general based on a number of factors including its own direct and indirect experiences (e.g., observations or deduction), on other agents' experiences, on an agents' reputation (Aberer & Despotovic, 2001; Mui et al., 2002). In general, trust is not transitive (i.e. recommendations can not be passed on), and trust is context dependent and dynamic. It has also been noted that trust and distrust are most often reciprocal by nature (Lawson, 1997; Falcone & Castelfranchi, 2001).

In some situations an agent will trust an agent fully, in others it may not. That an individual's reputation (Baya & Leifer, 1996; Lang et al., 2002) influences design processes is recognized. It clearly also influences the trust other designers have in other individual agents.

The trust agents have in each other changes over time. Agents continually update beliefs in other agents (Birk 2001; Witkowski et al., 2001; Beth et al., 1994; Barber & Kim, 2001), themselves and their environment. They continually need to deal with the "trust dilemma": the trade-off between positive or negative results of trusting another agent versus positive or negative results of not trusting another agent. Trust and risk are co-related concepts; trust is used to rationalise decisions involving risk (e.g., McKnight et al., 2002).

In small environments direct interaction-derived reputation-based trust mechanisms may suffice (Birk, 2000; Witkowski et al., 2001) to model the role trust plays. In open, dynamic environments in which incentives may differ it is questionable whether agents will always be truthful with respect to the information they provide about other agents (Jurca & Faltings, 2002; Abdul-Rahman & Hailes, 2000; Schillo et al., 1999; Beth et al., 1994).

### **2.3 TRUST MODELS**

Existing models of trust express: acquisition and representation of trust (e.g., Castelfranchi & Falcone, 2000), communication of trust (e.g., Aberer & Despotovis, 2001), and reasoning about trust (e.g., Marsh, 1994). Different trust models have been devised for different purposes, a number of which are discussed in this section.

Our knowledge-level approach to modelling individuals' involvement in distributed design processes yields insight in their aggregate behaviour. To explore the role of trust, trust models are sought that support this approach. This implies that applicable trust models need to be knowledge-level models: they need to be defined in terms of intentionalistic notions such as beliefs, desires, and intentions. Explicit representation of the multi-faceted nature of trust is also of importance for reasoning with and about trust. To facilitate experimentation, applicable trust models also need to be operationalisable.

Trust models such as Bell-LaPadula (Bell & LaPadula, 1973) distinguish different levels of trust and the relations between them. Such models require explicit knowledge of the levels of trust within a domain, and the role of individual agents in this hierarchy. The Bell-LaPadula model is designed for the military, in which such roles are clearly distinguished. Such relations are less easily defined for open distributed systems in which large numbers of agents operate. The Bell-LaPadula model is symbolic and not a knowledge-level model; trust is not expressed as a single value. It is operationalised: a number of implementations of Bell-LaPadula exist.

Marsh (1994) models trust in a simple but expressive way. In his model an agent's trust in another agent is based on three types of trust: basic, general and situational trust. A value is assigned to each type (a continuous value between -1 inclusive and 1 exclusive). Basic trust expresses the disposition of an agent towards trust in general. General trust, the trust of an agent in a specific other agent unrelated to situation or context, may be influenced by the value of basic trust. Situational trust, the trust of an agent in a specific other agent in a specific situation or context, is based on an estimation of general trust combined with the utility and importance of the situation. The threshold used to determine whether an agent trusts another

agent sufficiently to co-operate with the other agent, is based on perceived risk, competence and importance. A decision to co-operate is made when the situational trust is higher than the threshold. Marsh's trust model is to some extent a knowledge-level model and different types of trust are distinguished, facilitating reasoning about and with trust. It is, however, not operationalised.

TrustBuilder (Winslett et al., 2002) is a trust management system. Its purpose is to negotiate which resources agents may access, based on both specific agent's and Trustbuilder systems' policies and credentials. Both agents and the TrustBuilder system have credentials: e.g., proof of membership of an organization, a credit card number or specific authorisation certificates. Both also have policies describing the parties to whom they will disclose those credentials. During negotiation a TrustBuilder system and an agent exchange policies and credentials until the TrustBuilder system is satisfied, or until one of the parties is unable or unwilling to provide the requested credentials. TrustBuilder does not determine how credentials (the basis for trust) are acquired. Trustbuilder is symbolic and not a knowledge-level model. The trust valuation does not reflect its multi-faceted basis. It is operationalised.

Aberer and Despotovic (2001) combine a trust model with a peer-to-peer (P2P) trust storage model, storing complaints of agents about interactions with other agents in a global distributed model. Trust values are either 0 or 1, i.e., dishonest or honest. An agent that is considering interaction with another agent asks its neighbours (peers) if they trust the other agents: in this model this translates to having knowledge of complaints and/or complaints about complaining agents. If the number of complaints found does not exceed an agent's threshold for acceptance it may decide to initiate interaction. This simple model may be suitable for straightforward applications such as trading communities (e.g. Ibazar). This trust storage model is a knowledge-level model and explicitly involves the basis for trust, facilitating reasoning about and with trust. It is also operationalised. However, its assumption that trust is based on complaints is not easily (or generally) applied to a distributed design setting: more aspects are of importance than conflicts (which can be considered as a kind of complaints).

Ramchurn et al. (2003) describe a trust model based on confidence and reputation in the context of negotiation and contracts. Confidence information is acquired on the basis of an agent's personal experiences, reputation information is gathered from other agents. Trust in a specific context has a value between 0 and 1, and is computed by weighing confidence information stronger than reputation information. This is a knowledge-level model, in which confidence and reputation information are combined into a single trust value. It is not multi-faceted. It is unclear whether this model has been operationalised.

Castelfranchi and Falcone (2000) describe a more elaborate trust model. They distinguish seven types of beliefs related to trust. These types of beliefs can be described as follows:

- *Competence belief*: belief that the other agent has the abilities to do the tasks.
- *Disposition belief*: belief that the other agent is inclined do what it says it will do.
- *Dependence belief*: belief that it is better to rely/depend on the other agent than to approach a task without the other agent.
- *Fulfilment belief*: belief that the goal will be achieved due to the other agent's contribution.
- *Willingness belief*: belief that the other agent has decided and intends to do an action (to achieve the goal).
- *Persistence belief*: belief that the other agent is stable in its intentions (related to reliability).
- *Self-confidence belief*: belief that the other agent knows that it can do an action.

Based on these beliefs and their subjective certainty a degree of trust is calculated. The decision to delegate, rely or bet on an other agent is based on a comparison of the degree of trust with the risk factor and possibly a self-trust estimation. This trust model is a knowledge-

level model, is operationalised, and combines the values of its seven beliefs into a single value of trust.

## 2.4 OUR MODEL

The seven beliefs related to trust, distinguished by Castelfranchi and Falcone (1998; 2000), described above, provide a means to structure trust encountered in distributed design. However, the computational model Castelfranchi and Falcone presented (1998), combines these beliefs in a single-valued expression of trust, thereby reducing the options for agents to explicitly reason about different aspects of trust. A different computational model is needed to model trust relations in distributed design.

The model designed for this purpose is based on the seven beliefs, but uses a different mechanism to express trust. A threshold (vector) is used to express the threshold values for each of the seven beliefs: the values needed for one agent to trust another in a given context. These thresholds are determined for each agent with respect to each other agent. The thresholds change over time, as do the beliefs agents have in each other. Comparing agent A's threshold values with the beliefs agent A has in agent B, provides enough information to determine to which extent agent A trusts agent B. (Note that agent A may be equal to agent B). The use of thresholds is similar to Marsh's (1994) approach: our beliefs are similar to Marsh's situational trust, and our threshold is compared with our beliefs, akin to comparing Marsh's threshold with Marsh's situational trust.

The proposed model uses simple values to express knowledge-level beliefs and thresholds: discrete values from -2 to +2 have been used to facilitate computation and to have a sufficiently explicit representation of 'degrees of belief'. Knowledge-level valuations are mapped into a discrete domain, on the assumption that five-valued predicates are sufficient to elicit information from humans (e.g., the Likert (1932) scale).

The mechanism(s) with which the valuations of the aforementioned beliefs are determined are not discussed in this paper. These beliefs can, e.g., be based on agents' reputations, based on personal experiences or experiences by other agents (cf. Aberer & Despotovic, 2001). If beliefs have continuous values then they need to be mapped to discrete values when used in this trust model, requiring explicit choices in the assessment of the beliefs in agents in specific contexts. This is not discussed in this paper, as in our experiment, a knowledge-level analysis yielded five values per belief.

As stated above the belief of an agent in another specific agent in a specific context is expressed as a vector of seven values corresponding to the seven beliefs distinguished by Castelfranchi and Falcone:

```
Current_Belief_in( agent, context, < 2, 1, 1, 2, 2, 2, 1 > )
```

The threshold expressing the levels of belief required for each individual type of belief in a specific context, is also expressed as a vector of seven values: the values express the minimum value for which each specific belief is deemed to be acceptable (in combination with the other six).

```
Threshold( context, < 1, 2, 1, 2, 2, 2, 2 > )
```

Evaluating the beliefs about agents with respect to a threshold corresponds to calculating the difference between the minimally required values and the actual belief values in a specific context, yielding discrete values between -4 and +4, inclusive. These extreme values arise when, e.g. the threshold specifies a value of +2, yet the belief-value is -2, or vice versa. The evaluation has negative values for threshold violations, zeros for exact matches, and positive values for exceptional satisfactions:

```
Evaluation( agent, context, < 1, -1, 0, 0, 0, 0, -1 > )
```

Comparing evaluations of agents and choosing the best alternative may involve more complex algorithms. For the domain of application discussed in the next section, namely distributed website design, a simple heuristic is used. This heuristic is based on counting the number of violations, exact matches, and exceptional satisfactions. This also makes it possible to choose between 'worst' candidates, a feature which may be relevant in situations in which work simply has to be performed. In the case of a tie, a random choice is made. The algorithm is shown in the pseudocode below:

```

Let v be the set of violaters
Let m be the set of minimal compliers
Let s be the set of superb compliers

For each evaluation e:
    add to v if e has one or more negative elements.
    add to s if e has one or more positive elements and no negative elements.
    add to m if e has no positive elements and no negative elements.

If s has one element,
    return that element as the best evaluation
else if s has more than one evaluation
    return the evaluation with most positive elements
else if m has one evaluation
    return that evaluation
else if m has more than one evaluation
    return an evaluation
else
    return the evaluation in s with the least number of negative elements

```

*Figure 2: Pseudocode for evaluation selection. The next section demonstrates the use of this trust model.*

### 3. DISTRIBUTED WEBSITE DESIGN

Distributed website design in which a number of team members each with their own expertise and experience, collectively design and build a website, is an example of a distributed design process in which trust plays an important role. In this section, the trust model presented above is used to analyse two specific design processes<sup>1</sup> and to evaluate the trust model presented above. A computational simulation of these processes demonstrates how trust and trust thresholds influence the design process' results. Section 3.1 describes the website design application. Section 3.2 depicts the specific beliefs and trust thresholds involved for task delegation. Section 3.3 illustrates the effect of different trust thresholds.

#### 3.1 DESIGN APPLICATION

Two specific distributed website design projects are described in this section: one for the design of a website for a non-profit organisation, the other is the design of a project for a company<sup>1</sup>. The same team is responsible for both projects and the design process is similar.

In general, the websites this design team produces include separate sections for: news, addresses and an interactive forum. The non-profit organisation requires a website to promote Physics and is aimed at high school students. The website is to contain information about the field of Physics, news items on progress in Science, explanatory articles, example exercises, addresses of professors, support and ideas for school projects, and information about BSc and MSc Physics programmes. This customer's emphasis is on the functionality of the website: on the information to be communicated and shared, and to a lesser extent on "look & feel". An extended example of such a website can be found at <http://www.natuurkunde.nl/> (in Dutch).

The other site is for a company. Its overall goal is to support and attract clientele. The website is to contain information about the products, addresses of dealers and repair shops, and discussion about technologies. This customer's emphasis is more on the look and feel, than on the functionality provided by the website.

The design team includes graphical designers, logical designers, code designers (programmers), and HTML designers. A website consists of models for its graphical and logical layout, and (executable) specifications of its functionality and HTML. Briefly summarised, graphical designers design the overall "look & feel" of the website, logical designers design the "flow": the sequences of pages a client encounters, code designers write functional code, and HTML designers implement the layout in HTML pages.

Usually, the graphical and logical designers interact with a customer, and subsequently delegate work to the code designers and HTML designers. No strict hierarchy is enforced and each participant is able to delegate work to other participants. In all cases, the participant to whom a job has been delegated reports back to the participant who assigned the job. Each designer usually adds his or her own requirements and results ((partial) designs) to the overall project information, maintained by a version control system (e.g., CVS).

Work delegated to a designer (by the customer or by other designers) includes (partial) website descriptions, specifications/requirements (possibly qualified), and design project goals. The designer reports include information on partial models of the website, specifications (requirements) of website models, the extent to which specifications have been fulfilled and possible problems encountered, including conflicting specifications (assessments of requirements with respect to website descriptions), and whether project goals have been fulfilled, e.g., finishing work on time with the expected quality.

### 3.2 ROLE OF TRUST IN WEBSITE DESIGN

The role the trust model, described in Section 2.3, plays in website design, is illustrated in this section for task delegation within the design projects described above. Beliefs about other designers are evaluated with a trust threshold for different contexts. The resulting analysis was implemented in a computational simulation, using these beliefs and trust thresholds. The trace can be found at [http://www.iids.org/research/distributed\\_design/aiedam2004/](http://www.iids.org/research/distributed_design/aiedam2004/). The agents in this trace represent their human counterparts. Figure 3 shows part of this trace, in which the graphical designer agent deliberates about delegating work to one of the two HTML designer agents.

```

"Graphic Designer" (t=6)
input:
  Requirement Qualification Set:
  Design Process Evaluations:
    dpe-gd-1(fin-gd-1(dsgn-gd-1(addressbook)))
  Design Process Objectives:
  Assessment:
  DesignObjectDescription:
reasoning:
  Checking incoming customer requirements
  Checking for missing designs ****
  Designing [news]
    Delegating news [1, 2, 1, 2, 2, 2, 2] (t=1)
      Current_Belief in [HTMLer 1] = [1, 1, 1, 1, 1, 2, 2]
      Current_Belief in [HTMLer 2] = [2, 1, 1, 2, 2, 2, 1]
      [HTMLer 1] [news] <0, -1, 0, -1, -1, 0, 0> (3, 4, 0)
      [HTMLer 2] [news] <1, -1, 0, 0, 0, 0, -1> (2, 4, 1)
    Selected [HTMLer 2]
  Checking for done state
output:
  Requirement Qualification Set:
    rqs-gd-5(html-gd-2(news))
  Design Process Evaluations:
    dpe-gd-2(fin-gd-2(dsgn-gd-2(news)))

```

```

Design Process Objectives:
    dpo-gd-2 (udod-gd-2 (dod-gd-2))
Assessment:
DesignObjectDescription:
dod-gd-2 (dsgn-gd-2 (news))

```

**Figure 3:** Excerpt from trace illustrating the graphical designer agent’s reasoning about his beliefs and trust in the HTML designer agents.

In this example, the graphical designer needs one of the two HTML designers to implement the layout of the news section of the website. The graphical designer’s beliefs about the two HTML designers are expressed as the graphical designer’s agent’s context-specific valuations of its belief predicates concerning the fitness of two agents. In this context the graphical designer needs an HTML designer who is capable of implementing the graphical design requirements of the news section of the website in one day. These beliefs are expressed as follows for the graphical designer agent:

```

Current_Belief_in( htmler_1, newsDesign, < 1, 1, 1, 1, 1, 2, 2 > )
Current_Belief_in( htmler_2, newsDesign, < 2, 1, 1, 2, 2, 2, 1 > )

```

The semantics of the order of the beliefs (and trust-thresholds) in the notation used, is:

```

< Competence, Disposition, Dependence, Fulfilment, Willingness, Persistence, Self-
confidence >

```

The interpretation of the belief about the first HTML designer agent is given below; this agent has

- Competence 1: which indicates that he/she has all the competencies needed.
- Disposition 1: usually does what he/she promises.
- Dependence 1: is a better alternative than using myself.
- Fulfilment 1: usually designs HTML satisfying requirements on time.
- Willingness 1: is normally willing to work on this task.
- Persistence 2: is very reliable.
- Self-confidence 2: is well aware of his/her layouting abilities.

The graphical designer agent employs the following trust threshold for this context:

```

Threshold( newsDesign, < 1, 2, 1, 2, 2, 2, 2 > )

```

The rationale for this threshold is that the graphical designer wants to be convinced that the designer who is to do the HTML work has a high likelihood of success for this task. The knowledge-level concepts with which this is modelled and on which these values are based, are as follows:

- Competence of  $\geq +1$ : is reasonably competent for this context.
- Disposition of  $\geq +2$ : always does what he/she agrees upon.
- Dependence of  $\geq +1$ : is a better alternative than myself.
- Fulfilment of  $\geq +2$ : will definitely successfully complete the task.
- Willingness of  $\geq +2$ : is really willing to take on the task.
- Persistence of  $\geq +2$ : is very reliable.
- Self-confidence of  $\geq +2$ : has full self-confidence in him/herself for this task.

In this example the following evaluation of the two HTML designer agents is obtained

```

Evaluation( HTMLer_1, newsDesign, < 0, -1, 0, -1, -1, 0, 0 > )
Evaluation( HTMLer_2, newsDesign, < +1, -1, 0, 0, 0, 0, -1 > )

```

In other areas of application comparing evaluations of the parties involved and choosing the best alternative may involve more complex algorithms. In this domain of application, a simple heuristic is used, based on counting the number of violations, exact matches, and exceptional satisfactions. In the running example, this amounts to:

```
Evaluation-Summary( htmler-1, newsDesign, < 3, 4, 0 > )  
Evaluation-Summary( htmler-2, newsDesign, < 2, 4, 1 > )
```

On the basis of this algorithm no satisfactory candidates are found (as the model of both HTML designers violates at least two threshold values), so the most promising HTML designer has to be chosen. HTML designer 2 is preferable over HTML designer 1 as the former has the most exact matches and exceptional satisfactions.

### 3.3 ROLE OF TRUST IN WORK DELEGATION

The design of the news section of the website entails collaboration between designers of HTML and designers of code. The code designers provide the code needed to implement the functionality and flow depicted in HTML pages, based on the logical requirements with which they are provided. This design of code and HTML for the news section is used to illustrate the effect of changing trust threshold on results of the design process.

#### *First project*

The following situation occurs in the design of the website for the non-profit organisation (the Physics website). Code designer 2 is responsible for adding functionality to the HTML which the HTML designer has provided for the news section of the website. To this end, the code designer 2 formulates requirements concerning the data items to be stored for a news item, e.g., title, abstract, author, section titles, paragraphs, images, movies, links, etc. A number of these requirements are violated in the HTML for the news section, provided by HTML designer 2: a conflict occurs. Code designer agent 2 has a number of options to resolve this conflict:

1. the code designer doesn't change his own requirements, and changes the HTML for the news section accordingly without consulting a HTML designer agent; possibly resulting in sub-optimal HTML.
2. the code designer asks HTML designer 2 to change the HTML for the news section; this may take some time but results in suitable HTML.
3. alternatively, the code designer may ask HTML designer 1 for help; this may take more time, as this HTML designer was unfamiliar with this part of the website.
4. the code designer works in parallel with a HTML designer; this implies that the code designer quickly fixes the HTML for his own purposes, and later integrates the properly fixed HTML from the HTML designer.
5. the code designer drops his own requirements.

The last option conflicts with the logical design requirements which state the need for elaborate news items, including more possibilities for images, text-layout and external links than supported by the (assumed) design requirements of HTML designer a 2. The fourth option, involving parallelism, was not favoured by the code designer as quickly fixing HTML may involve as much work as doing it properly, and has the drawback that other conflicts may arise.

Code designer 2's beliefs are modelled according to the trust model, represented by code designer 2's agent. The beliefs about HTML designer 2 show the code designer's bias about this HTML designer, who originally designed the HTML for the news section. The code designer's agent also explicitly models beliefs about the code designer himself, as he is a possible candidate for work in this context.

```
format: < Compet, Dispos, Depend, Fulfil, Willing, Persist, Self-conf >  
Belief( progger-2, fixNewsHtml, < 1, 2, 2, 2, 2, 2, 1 > )
```

```
Belief( htmler-1, fixNewsHtml, < 1, 1, 0, 1, 1, 2, 2 > )
```

```
Belief( htmler-2, fixNewsHtml, < 2, 2, 1, 2, 1, 2, 2 > )
```

The code designer agent has a trust threshold based on the code designer's principle that making progress is more important than realising graphical quality. This results in the following trust threshold, acquired during our knowledge-level analysis, in which willingness and persistence are shown to be more important than other aspects:

```
Threshold_1( fixNewsHtml, < 1, 1, 1, 1, 2, 2, 1 > )
```

The following evaluation of the beliefs about the three parties is obtained:

```
Evaluation( progger-2, fixNewsHtml, < 0, +1, +1, +1, 0, 0, 0 > ) = 0,4,3
```

```
Evaluation( HTMLer_1, fixNewsHtml, < 0, 0, -1, 0, -1, 0, +1 > ) = 2,4,1
```

```
Evaluation( HTMLer_2, fixNewsHtml, < +1, +1, 0, +1, -1, 0, +1 > ) = 1,2,4
```

In this situation, the models of both HTML designers have at least one violation, and the model of the code designer himself has no violations at all. As a result, the code designer is shown to be the best candidate for fixing the HTML. In other words, the code designer agent 2 decides to choose option 1: to have the code designer fix the HTML himself.

### *Second project*

The same situation occurs in the second design project considered: the website for company, for which look & feel are more important. The same conflict between code designers and HTML designers occurs. In this case the code designer believes the quality of graphics is more important than making progress. His agent is therefore implemented with a trust threshold based on this. This trust threshold, acquired during our knowledge-level analysis, expresses the fact that competence, fulfilment, reliability and self-confidence are more important:

```
Threshold_2( fixNewsHtml, < 2, 1, 1, 2, 1, 2, 2 > )
```

The following evaluation of the beliefs about the three parties is obtained:

```
Evaluation( progger-2, fixNewsHtml, < -1, +1, +1, 0, +1, 0, -1 > ) = 2,2,3
```

```
Evaluation( HTMLer_1, fixNewsHtml, < -1, 0, -1, -1, 0, 0, 0 > ) = 3,4,0
```

```
Evaluation( HTMLer_2, fixNewsHtml, < 0, +1, 0, 0, 0, 0, 0 > ) = 0,6,1
```

In this situation, the code designer agent and the model of HTML designer 1 have at least two violations, while the model of HTML designer 2 has no violations at all. HTML designer 2 is chosen as the candidate for fixing the HTML. As a result code designer 2 has time to do other work, before resuming work on implementing code for the HTML of the news section.

### *Role of trust*

In both situations described above, the code designer is entrusted with a conflict with respect to the news section. He needs to extend the information stored for news items; the website design provided by the HTML designer is incomplete. The code designer's approach to resolve this conflict is based on characteristics of the current customer, as shown above. The resulting HTML for the news items is of a better quality in the second case than in the first case (mostly noticeable in details).

The two cases illustrate the role of trust in deliberating delegation of work, and subsequent effects on a design process. The combination of explicitly modelling seven beliefs for a specific agent in a specific context and employing a trust threshold for each of the seven beliefs facilitates the explication of trust in both cases. For this example the human expert's notions of trust were fairly easily captured in the trust model.

## 4. DISCUSSION

The role of trust in distributed design is the main topic of this paper. Human designers most often unconsciously assess the trust they have in other members of a design team, judging the value of their input accordingly (Milewski & Lewis, 1997). Automated support of distributed design requires an understanding of these processes. Fully automated distributed design requires computational models of distributed design including explication of trust relations. As yet it is infeasible to automatically assign trust levels to agents (see Falcone et al., 2001) because of the dynamic nature of trust. These relations will, however, need to be understood if automated support is to play a significant role in distributed design.

The trust model presented in this paper is based on Castelfranchi and Falcone's model (2000), in which seven beliefs are related to trust. Castelfranchi and Falcone (1998) compute a single value of trust on the basis of seven beliefs; the trust model in this paper uses a more expressive variant involving seven separate values for both trust thresholds and trust evaluation.

The example used to illustrate the potential of the model is that of task delegation within a design team. The reasoning processes involved are similar to the reasoning processes involved in determining the risk in taking a (co-operative) action (Griffiths & Luck, 1999). Griffiths and Luck specify trust related to the notion of general trust (Marsh, 1994), and explicitly omit situational trust for being too computationally expensive. The example in Section 3 illustrates the role of trust in task delegation within a closed environment. Different trust valuations results in different results. This example is based on the experiences of the human code designer and his rationalisation of beliefs and trust thresholds. The design results differ in both quality of the website (albeit in details), different sets of (qualified) requirements and different usage of time for involved designers.

Current research focuses on a different domain of application: the role of trust in distributed system location management. The simple model of trust reported in (Brazier & Wijngaards, 2002) is being replaced by the model reported in this paper. AgentScape (Wijngaards et al., 2002) is a world-wide scalable distributed agent platform. Management of AgentScape sites (i.e. locations within AgentScape) is a fully automated distributed configuration problem, an example of automated distributed design (Brazier & Wijngaards, 2002). An essential element in this design problem is its dynamic and open nature: the environment changes and the local configuration adapts.

Open environments place additional requirements on a trust model. The most obvious is the need to determine ways to acquire and adapt the values for the beliefs distinguished. Current experience in the above mentioned domain is promising. Most beliefs can be identified. Additional factors, however, play a role (e.g., factors related to an agent's environment such as security). TrustBuilder should work well within open environments, as does the approach based on distributed trust in open multi-agent systems, based on certificates, described by Mass and Shehory (2001) (as an extension of work done by Wong and Sycara (2000)). Mass and Shehory's approach allows agents to establish trust among themselves and update this trust when necessary without necessarily identifying themselves explicitly, basically a reputation-based scheme. Comparing incentives is less easily achieved in open environments as most individual incentives are not easily compared.

Our research on the role of trust in distributed design is just a beginning, more research is clearly needed.

## ACKNOWLEDGEMENTS

The authors express their gratitude to Cristiano Castelfranchi for his comments on this manuscript, to Us Media (<http://www.usmedia.nl/>) for providing material for the website design example and to Stichting NLnet, <http://www.nlnet.nl/> for their support.

## REFERENCES

- Abdul-Rahman, A., & Hailes, S. (2000). Supporting Trust in Virtual Communities. *Proceedings 33rd Hawaii International Conference on System Sciences*, volume 6, pp. 6007. IEEE Press.
- Aberer, K., & Despotovic, Z. (2001). Managing Trust in a Peer-2-Peer Information System. *Proceedings of the 10th International Conference on Information and Knowledge Management*, pp. 310-317. ACM Press.
- Adams, R.S., Turns, J. & Atman, F.J. (2003). Educating effective engineering designers: the role of reflective practice, *Design Studies*, 24(3), 275-294.
- Anumba, C.J., Ugwu, O.O., Newnham, L. & Thorpe, A. (2002). Collaborative design of structures using intelligent agents. *Automation in Construction*, 11(1), 89-103.
- Atman, C.J., Chimka, J.R., Bursic, K.M. & Nachtmann, H.L. (1999). A comparison of freshman and senior engineering design processes. *Design Studies*, 20(2), 131-152.
- Barber, K.S., & Kim, J. (2001). Belief Revision Process Based on Trust: Agents Evaluating Reputation of Information Sources. *Trust in Cyber-societies, Integrating the Human and Artificial Perspectives* (Falcone, R., Singh, M.P., & Tan, Y-H., Eds.), Lecture Notes in Computer Science, 2246, pp. 73-82. Springer.
- Baya, V., & Leifer, L.J. (1996). Understanding Information Management in Conceptual Design. *Analysing Design Activity* (Cross, N., Christiaans, H., & Dorst, K., Eds.), pp. 151-168. Chichester: John Wiley & Sons.
- Bell, D.E., & LaPadula, L.J. (1973). *Secure Computer Systems: Mathematical Foundations and Model*. The MITRE Corporation, report MTR 2547 v2.
- Beth, T., Borcharding, M., & Klein, B. (1994). Valuation of Trust in Open Networks. *Proc. 3rd European Symposium on Research in Computer Security*, pp. 3-18.
- Birk, A. (2000). A Boosting Cooperation by Evolving Trust. *Applied AI*, 14, 769-784.
- Birk, A. (2001). Learning to trust. *Trust in cyber-societies: integrating the human and artificial perspectives* (R. Falcone, R., Singh, M., & Tan, Y-H., Eds.), Lecture Notes in Computer Science, 2246, pp. 133-144. Springer.
- Boujut, J-F. & Laureillard, P. (2002). A co-operation framework for product-process integration in engineering design. *Design Studies*, 23(6), 497-513.
- Brazier, F.M.T., & Wijngaards, N.J.E. (2002). The Role of Trust in Automated Distributed Design. *Proceedings of the Workshop on Agents in Design* (Gero, J.S., & Brazier, F.M.T., Eds.), pp. 71-84. Australia: University of Sydney.
- Brazier, F.M.T., Jonker, C.M., Treur, J., & Wijngaards, N.J.E. (2001b). Compositional Design of a Generic Design Agent. *Design Studies*, 22(5), 439-471.
- Brazier, F.M.T., Moshkina, L.V., & Wijngaards, N.J.E. (2001a). Knowledge level model of an individual designer agent in collaborative distributed design. *Journal of Artificial Intelligence in Engineering*, 15, 137-152.
- Busby, J.S. (2001). Error and distributed cognition in design. *Design Studies*, 22(3), 233-254.
- Camp, L.J. (2003). Design for Trust. *Trust, Reputation and Security: Theories and Practice* (Falcone, R., Barber, S., Korba, L. & Singh, M., Eds.), volume 2631, pp. 15-29. Berlin: Springer-Verlag.
- Campbell, M.I., Cagan, J., & Kotovsky, K. (1998). A-Design: theory and implementation of an adaptive, agent-based method of conceptual design. *Artificial Intelligence in Design '98* (Gero, J.S., & Sudweeks, F., Eds.), pp. 579-598. Dordrecht: Kluwer Academic Publishers.
- Castelfranchi, C., & Falcone, R. (1998). Principles of Trust for MAS: Cognitive Anatomy, Social Importance, and Quantification. *Third International Conference on Multi Agent Systems*, pp. 72-80.
- Castelfranchi, C., & Falcone, R. (2000). Trust is Much More Than Subjective Probability: Mental Components and Sources of Trust. *Proceedings 33rd Hawaii International Conference on System Sciences*, volume 6, pp. 6008. IEEE Press.
- Chao, K-M., Norman, P., Anana, R., & James, A. (2002). An Agent-Based Approach to Engineering Design. *Computers in Industry*, 48, 17-27.
- Coates, G., Duffy, A.H.B., Hills, W., & Whitfield, R.I. (2000). A Generic Coordination Approach Applied to a Manufacturing Environment. *Journal of Materials Processing Technology*, 107, 404-411.

- Craig, D.L. & Zimring, C. (2000). Supporting collaborative design groups as design communities. *Design Studies*, 21(2), 187-204.
- Cross, N., Christiaans, H., & Dorst, K. (1996). *Analysing Design Activity*. Chichester, England: John Wiley & Sons.
- Denton, H.G. (1996). Developing Design Teamworking Capability: some planning factors emerging from a survey of engineering design courses. *Proceedings of the International Conference on Design and Technology Educational Research and Curriculum Development* (Smith, J.S., Eds.), pp. 112-117. Loughborough University, IDATER 96.
- Dwarakanath, S., & Blessing, L. (1996). Ingredients of the Design Process: a Comparison between Group and Individual Work. *Analysing Design Activity* (Cross, N., Christiaans, H., & Dros, K., Eds.), pp. 93-116. Chichester: John Wiley & Sons.
- Falcone, R., & Castelfranchi, C. (2001). The Socio-cognitive Dynamics of Trust: Does Trust Create Trust? *Trust in Cyber-societies, Integrating the Human and Artificial Perspectives* (Falcone, R., Singh, M.P., & Tan, Y-H., Eds.), Lecture Notes in Computer Science, 2246, pp. 55-72. Springer.
- Falcone, R., Singh, M.P., & Tan, Y-H. (2001). Introduction: Bringing Together Humans and Artificial Agents in Cyber-Societies: A New Field of Trust Research. *Trust in Cyber-societies, Integrating the Human and Artificial Perspectives* (Falcone, R., Singh, M.P., & Tan, Y-H., Eds.), Lecture Notes in Computer Science, 2246, pp. 1-8. Springer.
- Gambetta, D. (2000). Can We Trust Trust. *Trust: Making and Breaking Cooperative Relations* (Gambetta, D., Eds.), pp. 213-237. Department of Sociology, University of Oxford.
- Gero, J.S. (1998). Conceptual Designing as a Sequence of Situated Acts. *Artificial Intelligence in Structural Engineering* (Smith, I., Eds.), pp. 165-177. Berlin: Springer.
- Gero, J.S., & Kannengiesser, U. (2002). The situated function - behaviour - structure framework. *Artificial Intelligence in Design '02* (Gero, J.S., Eds.), pp. 89-104. Kluwer Academic Publishers.
- Greco, D.L., & Brown, D.C. (1996). Learning by single function agents during spring design. *Artificial Intelligence in Design '96* (Gero, J.S., & Sudweeks, F., Eds.), pp. 409-428. Dordrecht: Kluwer Academic Publishers.
- Griffiths, N., & Luck, M. (1999). Cooperative Plan Selection Through Trust. *Multi-Agent System Engineering - Proceedings of the Ninth European Workshop on Modelling Autonomous Agents in a Multi-Agent World* (Garijo, F.J., & Boman, M., Eds.), Lecture Notes in Computer Science, 1647, pp. 162-174. Berlin: Springer-Verlag.
- Hertzum, M. (2002). The importance of trust in software engineers' assessment and choice of information sources. *Information and Organization*, 12, 1-18.
- Hertzum, M. & Mark Pejtersen, A. (2000). The information-seeking practices of engineers: searching for documents as well as for people. *Information Processing & Management*, 36(5), 761-778.
- Hill, A.W., Dong, A. & Agogino, A.M. (2002). Towards computational tools for supporting the reflective team. *Artificial Intelligence in Design '02* (Gero, J.S., Eds.), pp. 305-325. Kluwer Academic Publishers.
- Holt, J.E. (1997). The designer's judgement. *Design Studies*, 18(1), 113-123.
- Jensen, D., Feland, J., Bowe, M. & Self, B. (2000). A 6-Hats Based Team Formation Strategy: Development and Comparison with an MBTI Based Approach. *Proceedings of the American Society for Engineering Education ASEE Annual Conference*. On CD, session 2425. St Louis.
- Jurca, R., & Faltings, B. (2002). Towards Incentive-Compatible Reputation Management. *Proceedings of the Workshop 'Deception, Fraud and Trust' of the Autonomous Agents Conference* (Falcone, R., Barber, S., Korba, L., & Singh, M., Eds.), pp. 92-100. ACM Press.
- Lang, S.Y.T., Dickinson, J., & Buchal, R.O. (2002). Cognitive Factors in Distributed Design. *Computers in Industry*, 48(1), 89-98.
- Lawson, B. (1997). *How Designers Think: The Design Process Demystified*, 3rd edition. Oxford: Architectural Press.
- Lees, B., Branki, C. & Aird, I. (2001). A framework for distributed agent-based engineering design support. *Automation in Construction*, 10(5), 631-637.
- Lewis, W.P. & Bonollo, E. (2002). An analysis of professional skills in design: implications for education and research. *Design Studies*, 23(4), 385-406.
- Likert, R. (1932). *A Technique for the Measurement of Attitudes*. New York: Archives of Psychology.
- Liu, H. & Frazer, J.H. (2002). Supporting evolution in a multi-agent cooperative design environment. *Advances in Engineering Software*, 33(6), 319-328.
- Luhmann, N. (2000). Familiarity, Confidence, Trust: Problems and Alternatives. *Trust: Making and Breaking Cooperative Relations* (Gambetta, D., Eds.), Chapter 6, pp. 94-107. Department of Sociology, University of

Oxford.

- Maher, M.L., Simoff, S., & Cicognani, A (2000). *Understanding Virtual Design Studios*. London: Springer-Verlag.
- Marsh, S.P. (1994). *Formalising Trust as a Computational Concept*. Department of Computing Science and Mathematics. University of Stirling.
- Mass, Y. & Shehory, O. (2001). Distributed Trust in Open Multi-agent System. *Trust in Cyber-societies, Integrating the Human and Artificial Perspectives* (Falcone, R., Singh, M.P., & Tan, Y-H., Eds.), Lecture Notes in Computer Science, 2246, pp. 111-132. Springer.
- McAlinden, L.P., Florida-James, B.O., Chao, K-M, Norman, P.W., Hills, W., & Smith, P. (1998). Information and knowledge sharing for distributed design agents. *Artificial Intelligence in Design '98* (Gero, J.S., & Sudweeks, F., Eds.), pp. 537-556. Dordrecht: Kluwer Academic Publishers.
- McKnight, D.H. & Chervany, N.L. (2001). Trust and Distrust Definitions: One Bite at a Time. *Trust in cyber-societies : integrating the human and artificial perspectives* (Falcone, R., Singh, M.P. and Tan, Y.-H., Eds.), Lecture Note in Computer Science, 2246, pp. 27-54, Springer.
- McKnight, D.H., Choudhury, V. & Kacmar, C. (2002). Developing and Validating Trust Measures for e-Commerce: An Integrative Typology. *Information Systems Research*, 13(3), 334-361.
- McKnight, D.H., Cummings, L.L. & Chervany, N.L. (1998). Initial Trust Formation in New Organizational Relationships. *Academy of Management Review*, 23(3), 473-490.
- Milewski, A.E. & Lewis, S.H. (1997). Delegating to software agents. *International Journal of Human-Computer Studies*, 46(4), 485-500.
- Mui, L., Mohtashemi, M., & Halberstadt, A. (2002). Notions of reputation in multi-agents systems: a review. *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pp. 280-287. ACM Press.
- Pynadath, D.V. & Tambe, M. (2002). Multiagent Teamwork: Analyzing the Optimality and Complexity of Key Theories and Models. *First Autonomous Agents and Multiagent Systems Conference (AAMAS02)*, pp. 873-880.
- Ramchurn, S.D., Sierra, C., Godo, L. & Jennings, N.R. (2003), A computational trust model for multi-agent interactions based on confidence and reputation. *Proc. 6th Int. Workshop of Deception, Fraud and Trust in Agent Societies*, Melbourne, Australia, pp. 69-75.
- Schillo, M., Funk, P., & Rovatsos, M. (1999). Who can you trust: Dealing with Deception. *Proceedings of the Workshop 'Deception, Fraud and Trust' of the Autonomous Agents Conference* (Castelfranchi, C., Tan, Y., Falcone, R. & Firozabad, B.S., Eds.). pp. 95-160. ACM Press.
- Schön, D.A. (1983). *The Reflective Practitioner: How Professionals Think in Action*. Basic Books. Perseus Books Group.
- Simoff, S.J. & Maher, M.L. (2000). Analysing participation in collaborative design environments. *Design Studies*, 21(2), 119-144.
- Sonnenwald, D.H. (1996). Communication roles that support collaboration during the design process. *Design Studies*, 17(3), 277-301.
- Tambe, M. (1997). Towards Flexible Teamwork. *Journal of Artificial Intelligence Research*, 7, 83-124.
- Valkenburg, R., & Dorst, K. (1998). The reflective practice of design teams. *Design Studies*, 19(3), 249-271.
- Wang, L., Shen, W., Xie, H., Neelamkavil, J., & Pardasani, A. (2002). Collaborative conceptual design - state of the art and future trends. *Computer Aided Design*, 34(13), 981-996.
- Wijngaards, N. J. E., Overeinder, B. J., Steen, M. van, & Brazier, F. M. T. (2002). Supporting Internet-Scale Multi-Agent Systems. *Data and Knowledge Engineering*, 41(2-3), 229-245.
- Wilson, J.L. & Shi, C. (1996). Co-ordination mechanisms for cooperative design. *Engineering Applications of Artificial Intelligence*, 9(4), 453-461.
- Winslett, M, Yu, T., Seamons, K.E., Hess, A., Jacobson, J. Jarvis, R., Smith, B., & Yu, L. (2002). Negotiating Trust on the Web. *IEEE Internet Computing*, 6(6), 30-37.
- Witkowski, M., Artikis, A. & Pitt, J. (2001). Experiments in Building Experiential Trust in a Society of Objective-Trust Based Agents. *Trust in Cyber-societies, Integrating the Human and Artificial Perspectives* (Falcone, R., Singh, M.P., & Tan, Y-H., Eds.), Lecture Notes in Computer Science, 2246, pp. 111-132. Springer.
- Wong, H.C. & Sycara, K.P. (2000). Adding Security and Trust to Multiagent Systems. *Applied Artificial Intelligence*, 14(9), 927-941.
- Wooldridge, M.J., & Jennings, N.R. (1995). Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10(2), 115-152
- Zha, X.F. (2002). A knowledge intensive multi-agent framework for cooperative/collaborative design modeling and decision support of assemblies. *Knowledge-Based Systems*, 15(8), 493-506.

## **BIOGRAPHIES**

### **prof.dr. Frances M.T. Brazier**

Frances Brazier is a full professor and heads the Intelligent Interactive Distributed Systems (IIDS) Group. Her research focuses on the design and development of support for large scale intelligent, interactive, distributed systems. This includes middleware (AgentScape), services (an agent factory and directory services), applications (distributed information retrieval and distributed design) but also consideration of legal implications (ALIAS). All in close collaboration with the IIDS group and others.

### **dr. Niek J.E. Wijngaards**

Niek Wijngaards is an assistant professor at the IIDS Group. He received his PhD in 1999, on the topic of self-modifying agent systems using a re-design process. His current research focusses on design processes to support large-scale heterogeneous adaptive multi-agent systems and analysis of legal implications of agent technology (ALIAS). One aspect is automated software configuration, including automated agent adaptation with agent factories and web-service configuration. Distributed design is studied as an example of a complex multi-agent system and as a service for the AgentScape agent operating system.

### **drs. Hidde M. Boonstra**

Hidde Boonstra is a Ph.D. student at the IIDS Group since 2003. He received his MSc in 2001 and has been actively involved in the development of commercial web-applications. His current research interests focus on security and trust in agent operating systems and agent systems. Distributed design is an example multi-agent application to test models and theories.