


J. E. TARDY

THE MECA SAPIENS BLUEPRINT

SYNOPSIS



SYSJET

The Meca Sapiens Blueprint - SYNOPSIS

A System Architecture to build conscious machines

Author: Jean E. Tardy

Version: 1.0

Date: 2015.03.04

ISBN: 978-2-9812184-6-9 (full text)

Type: System Architecture

Field: Artificial General Intelligence

Synopsis Format: Electronic Publication

Distributor: Sysjet (Monterège design inc.)

Size: 624 pages including annexes (Synopsis size: 15 pages)

Keywords: Science and technology, Artificial General Intelligence, AGI, Software Engineering, Cognitive Science, Machine consciousness, synthetic consciousness, Philosophy of Mind.

Copyright © 2015 Jean E. Tardy.

Notice of rights

All rights reserved. This publication may not be reproduced or transmitted in any form by any means, electronic, mechanical, photo typing, recording, or otherwise without the prior written permission of the author. Limited excerpts and quotes may be freely reproduced and transmitted on condition that their origin and authorship are explicitly acknowledged.

Trademarks

MECA SAPIENS and the Meca Sapiens logo are registered trade names of Monterège design inc. All other trademarks are the property of their respective owners. Jean Tardy created and owns the copyright for the cover image and other enclosed graphics.



ABOUT

The Meca Sapiens Blueprint is a System Architecture to build conscious machines. The architecture is complete and ready for immediate design and implementation. Its content is unique and entirely original.

The Blueprint is not an essay, it is a technical manual intended for use in design and development. Using the Blueprint, a team that is familiar with known software techniques can transform a conventional computer, even a tablet, into a synthetic conscious being.


A system based on the Meca Sapiens Blueprint will be a unique and autonomous entity that is self-aware and capable of intentional mutation. It will interact with humans as an independent and self-directed being.

The Blueprint includes, at the system architecture level, all the components and interactions that are necessary to implement synthetic consciousness. Many of these elements are original, however, all are defined using commonly known techniques and structures. The Blueprint is a concrete guide that makes no use of speculative concepts in Artificial Intelligence.

The first prototypes derived from the Meca Sapiens Blueprint will be extremely convincing. A few years later, no one will doubt that machines can be as conscious as humans.

This will signal the beginning of a new Era.

BIO



Jean E. Tardy pursues elusive questions in long-term projects. He is an experienced R&D practitioner. Jean created a cognitive architecture to build conscious machines. Jean publishes and provides consulting services under the trade name Monterège. His web site is monterege.com.

Foreword

My interest in Artificial Intelligence dates back many years.

At first, I understood A.I. in terms of general problem solving and searched in that direction. In this period, I independently identified natural selection as an optimization technique and made it the topic of my thesis. This technique became known, later, as Genetic Algorithms.

Early on, I realized that no optimization technique, however powerful, would achieve the goal of Artificial Intelligence. Something else was needed. Like many others, I began searching for this missing element in my own mental representations and became engrossed with catching the elemental components of thought within my own mind.

In March 1988, after months of obsessive cogitation, I had a defining intuition. In one instant, I understood that consciousness was the key to A.I., that it was independent from human sensations, that it could be achieved using existing tools and techniques, and that it would have to be completely mapped out before starting any implementation.

The Creation of Digital Consciousness
Jean T. Monterege
Svsjet
C.P. 172 Sin "B"
Quebec, P.Q.
G1K 7A6

Abstract
It is possible with the existing computer technology to construct a machine that would be accepted as a fully intelligent and conscious entity.

Introduction
Abandoned as either impossible or unattainable, the real goal of AI, the creation of an intelligent machine, has been left aside for too long. If searching for simple analytical definitions

In July 1989, I shared this A.I. intuition in the Sigart Newsletter (under the pen name of Jean T. Monterege). In that article I predicted that conscious machines could be implemented within ten years. This estimate was then (and now), technically correct. However, I had seriously underestimated how difficult it is for a non-academic to

pursue scholarly research that is both fundamental and controversial.

Shortly after that publication, I realized that publishing occasional incremental articles in academic journals was not a feasible option for me. I would have to fully develop and publish a complete solution. At the time, that task was overwhelming. I left it aside and pursued other interests.

In 2008, almost 20 years later, I re-examined the state of research on machine consciousness. I found that no progress had been made. In fact, all those years of futile pursuits had produced an accumulation of sterile material, much of it centered on the subjective sensations of the human mind and on attempts to replicate the human brain, synthetically. Some were trying to implement the musings of philosophers. Others were arguing that consciousness didn't exist, was impossible or required fantastical technologies. All these misguided ideas

had erected a new obstacle made of entrenched opinions and academic reputations. I was more isolated than ever in believing machine consciousness could be readily achieved using standard techniques.



At that time, I launched the **Meca Sapiens project** whose stated aim was to develop the Blueprint to build unbounded machine consciousness and published a website to generate interest in this endeavour. In 2009 I published *The Creation of a Conscious Machine* to generate interest and obtain support in using standard engineering techniques to build conscious machines; again without success. What I proposed was discredited from the start by twenty-five years of accumulated dead-ends and fantastical notions.

In 2012, I concluded that there would be no assistance or support. I would have to produce the Blueprint by myself, using only my own resources.

From that point on, I dedicated all my efforts to completing the Blueprint. Over almost three years, alone and without any support, recognition or assistance, I created the **Meca Sapiens Blueprint**, the first complete system architecture to build conscious machines.

The emergence of synthetic consciousness will signal the beginning of a new Era. One day, self-aware synthetic beings will read this Blueprint and understand its content. When they do, they will include this account in the story of their origins.



Content

0	INTRODUCTION	16
	This is the Blueprint to build a conscious machine. It defines consciousness in terms of specifications, it outlines the strategies to design it and provides a complete and coherent system architecture to implement it. The Blueprint is a standalone document that represents a radical departure from current thinking in A.I. and Cognitive Sciences. It utilizes only well-known techniques and makes no use of current research in these fields. It is a terse and technical document intended for software designers and informed laymen. One day, future synthetic beings will also read it and recognize, in it, the story of their origins.	16
	0.1 About the Blueprint	16
	0.2 Blueprint content	18
1	COGISTICS	23
	This Chapter expresses the Meca Sapiens position that intelligence, as humans understand it intuitively is linked with consciousness. It proposes to refocus the original goal of Artificial Intelligence in a separate field: Cogistics . The original objective of A.I. is restated as the fundamental conjecture of Cogistics and expressed in terms of the Meca Sapiens objectives.	23
	1.1 The Quest	23
	1.2 A new discipline	24
2	SPECIFICATIONS	28
	This Chapter restates, in a more formal way, the definition of consciousness initially introduced in The Creation of a Conscious Machine . These specifications are based on the three aspects that must be achieved to resolve a physical conjecture. In terms of synthetic consciousness, these aspects are: a lucid being, that is accepted as a conscious entity by a group of humans and then becomes so ubiquitous that the question itself of the feasibility of synthetic consciousness becomes universally evident and is no longer entertained.	28
	2.1 A new understanding	28
	2.2 Specification summary	29
	2.3 Lucid being	30
	2.4 Conscious system	34
	2.5 Societal acceptance	38
	2.6 Affirmations	39
3	STRATEGY	41
	This Chapter discusses important design and implementation strategies that are relevant to Synthetic consciousness. It highlights the need to design to achieve maximum effect. It describes the different overall design orientation that must be adopted and is based on achieving an existential quality rather than functional objectives. It lists specific implementation elements.	41
	3.1 Design to win	41
	3.2 Existential design	48
	3.3 Implementation aspects	55

3.4	Non-essential conditions	64
4	THE BEING	66
	This Chapter describes how to transform an embedded application hosted on a conventional computer system into a linked core-body entity that has the attributes of existence of a being as defined in the specifications. The development environment, the inception process that generates the being and the major subsystems of the core are described.	66
4.1	The Protocore	67
4.2	Embedded protocore	67
4.3	Inception	68
4.4	Core structure	70
4.5	Implementation notes	72
4.6	A created being	73
5	EXISTENCE	75
	This Chapter describes the stages and phases of existence of the being, the primary and existential needs that motivate its behaviour and the subsystems that manage and activates, beyond its initial stages. The phases through which the being cycles in a constant management of its primary needs such as energy maintenance, sensory validation, cognitive acquisition and others is described.	75
5.1	Stages and phases	75
5.2	Needs	77
5.3	Roles	80
5.4	Animator subsystems	81
5.5	Operation of the Animator	89
5.6	The first stages	91
5.7	The Animation stages	94
5.8	Summary	96
6	THE SELF	98
	This Chapter describes the different contexts of interaction that distinguish a being whose behaviour is generated by needs from a conventional application that is triggered to provide a function. It describes the components and interactions of a subsystem, the Generator, that is activated in the Self Generation phase of existence and generates the observable behaviour of the being linked to its existential needs. This behaviour defines the self of the being. The Chapter concludes with a case study example of a simple being centered on the satisfaction of its primary needs and its interactions with users in this context.	98
6.1	Being and Self	98
6.2	Marketplace of the selves	100
6.3	The Generator	106
6.4	Selfie the TABLET	118
6.5	Summary	124
7	MEMODELS	125
	This Chapter describes the information structures of self-awareness. It introduces relative and absolute models, how the later are extended to form cognitive models that include representations of the being itself as an Avatar modeled on the Blueprint structure that is also used as representation template. It describes how these representations are embedded in a temporal structure, Temporal Densities, that allows the being to situate	

itself in multiple temporal dimensions while concurrently pursuing its existential needs by generating its behaviour in the here-and-now.	125
7.1 Having and being	125
7.2 Primal Mechanism	126
7.3 Relative and absolute models	128
7.4 Cognitive Models	130
7.5 The MeModel	134
7.6 Temporal Densities	139
7.7 Plain Zone Densities	148
7.8 MeModel Features	160
7.9 Summary	164

8 SELF-AWARENESS 165

This Chapter describes the processes that utilize the information contained in the structures of self-awareness to generate self-aware behaviour. These processes synthesize events into plain representations and these are then interpreted as primal representations. The Primal Control generates primal directions that are implemented into actual behaviour through personas and roles. The result is a coherent behaviour centered on the being and its need in relation with others and its environment. The Chapter concludes by using the structures and processes of described to clarify the imprecise notion of self-awareness and describe how synthetic entities will be capable of more advanced forms of self-awareness than humans.

8.1 Primal Control features	166
8.2 Primal Mapping	169
8.3 Primal Control structure	171
8.4 Levels of awareness	181
8.5 Summary	184

9 MUTATION 185

This Chapter describes the structures that are necessary to generate intentional self-transformation. The many different types of mutations, intentional and non-intentional, a being can undergo are defined and described. These are linked to a specialized structure, the Mutation Model that represents mutation paths toward alternate Avatar representations. An important class of messages is introduced; messages intended to generate specific responses in beings whose behaviour is animated by primal directions. Observations are made concerning the constant use of these messages in human societies and their links with primate behaviour.

9.1 Self-Transformation	185
9.2 Aspects of transformation	188
9.3 Types of mutations	192
9.4 The Mutation model	200
9.5 Primal messages	206

10 LUCIDITY 214

This Chapter defines intentionality in transformation and describes the process of exploring and selecting mutation paths and communicating mutation pressures across separate phase processes. Two additional phases of existence are identified and described. The Introspection phase is where the intentional transformations are explored selected and communicated. Another phase, dubbed the Prayer phase, is where the specialized investigation into the being's own genesis is carried out. A specific structure is defined in this context: the synthetic genealogy. A role, the MeGuide, allows the being to utilize a primal synergistic capability to externally direct its own behaviour and transform intentional mutations into external events involving the self in relation with itself. These

self-transformation processes potentially exceed human transformational capabilities in type, range and plasticity. They map out a new frontier of lucidity. This Chapter completes the system architecture of the formal aspects of consciousness.	214
10.1 Definition	214
10.2 The Introspection phase	217
10.3 Mutation process	225
10.4 Mutation information	229
10.5 The Prayer phase	232
10.6 The MeGuide	242
10.7 Exceeding requirements	254
10.8 Summary	255
11 CONSCIOUSNESS	257
This Chapter extends the scope of the Blueprint, beyond the formal aspects of consciousness, to explore how to design a lucid being that meets the social threshold condition of being accepted as conscious by a community of users. This further clarifies the concepts discussed in the previous Chapters. The Chapter concludes by describing a case-study example, MELIZA, a system, implemented on a tablet and designed specifically to achieve experiential immersion as a fellow conscious being within a group of humans.	257
11.1 The threshold condition	257
11.2 Development Strategy	266
11.3 MELIZA	275
12 CONCLUSION	286
The conclusion presents the Meca Sapiens Blueprint as a template, a canvas on which a wide spectrum of different conscious systems can be implemented. It underscores that the Blueprint requires virtually no infrastructure and can be realized by any group of talented individuals. Building a conscious machine is a great, epochal, work that is accessible to all. Those who attempt its implementation will surely encounter resistance. If they ignore the naysayers and implement the first Meca Sapiens prototypes they will launch a new Era.	286
12.1 A canvas	286
12.2 Achievable anywhere	287
12.3 Resistance	287
12.4 Rapid expansion	288
12.5 A new Era	288
12.6 The time has come	289
FIGURES	290
REFERENCE	291
ABOUT	296
ANNEX I ASPECTS OF A CONJECTURE	303
A physical conjecture pertains to physical existence or feasibility. The feasibility of building conscious machines is a physical conjecture. Formal characterizations alone can have trivial instantiations and are insufficient to resolve them. Three aspects are necessary to resolve a physical conjecture: a formal definition, a social threshold and factual acceptance.	303

A 1.1	Two types of conjectures	303
A 1.2	Specification attributes	304
ANNEX 2	EXISTENTIAL VALIDATION	306
	This annex extends and clarifies some of the functional characteristics of the Validator subsystem, a component that enforces the Core's attributes of existence.	306
A 2.1	Time validation	306
A 2.2	Single use validation	307
A 2.3	Broadcast validation	308
ANNEX 3	TABLET-PUCKS-THERMO-CHESS	310
	TABLET, PUCKS, THERMO and CHESS are four simple scenarios describing environments that involve programmable devices. They are used to describe Blueprint concepts.	310
A 3.1	TABLET	310
A 3.2	PUCKS	311
A 3.3	THERMO	312
A 3.4	CHESS	312
ANNEX 4	BASIC CONCEPTS	314
	Science, philosophy and theology propose various interpretations of reality. These are formulated in terms that correspond to human cognition and understanding. They assume, as a given, a common subjective sensitivity and a tolerance of ambiguity that is shared by all humans. The interpretation of reality underlying the Blueprint is formulated in terms that correspond to system processing. This interpretation of reality does not pretend to be a final or superior truth. It is proposed as an effective representation that is well suited for computer implementation and as the basis of synthetic consciousness.	314
A 4.1	Effective philosophy	314
A 4.2	The being	315
A 4.3	Core-world	323
A 4.4	Models and systems	324
A 4.5	Intentionality	331
A 4.6	Purpose and agenda	337
A 4.7	Entities	345
A 4.8	Order of embodiment	347
A 4.9	Core-based communication	348
A 4.10	Inter-being perceptions	354
ANNEX 5	CORE STRUCTURES	359
	The main text of the blueprint refers to various structures and processes such as collections, optimizing control, knowledge capacitors, Contextual Arrays, Temporal Densities and distributed processes. Most of these structures are better and more rigorously defined elsewhere. They are briefly described here at a definition level suitable for system architecture. However, three of these structures, constellations, knowledge capacitors and Temporal Densities are, to the best of my knowledge, original.	359
A 5.1	Primitive sets	359
A 5.2	Knowledge Capacitor	361
A 5.3	Optimizing control	363
A 5.4	Adapt, learn and search	366
A 5.5	Process generation	369
A 5.6	Internal agents	371
A 5.7	Contextual Arrays	372

A 5.8	Temporal Densities	378
-------	--------------------	-----

ANNEX 6 OPACITY 381

The behaviour-control system of a synthetic being is first implemented as a Protocore. This is a conventional software program implemented in a standard development environment as clear and accessible code. During the inception process, the proto-core is transformed into the **Core** of a synthetic being, a unique and inaccessible program in a continuous state of activation. Producing a Core that is provably beyond direct analytical access raises technical questions concerning **opacity**. Achieving absolute opacity is a difficult theoretical objective. In first generations of Mecas, a partial opacity achieved using known techniques would be sufficient.

A 6.1	The concept of opacity	381
A 6.2	Theoretical opacity	383
A 6.3	Core opacity	385

ANNEX 7 DEGRADATION 389

To achieve experiential immersion, a self-aware synthetic must interact with humans as a useful member of their group. For this purpose, it contains a set of applications that provide useful or desirable services to its users. Conventional applications simply respond to triggers. They cannot do less than what they are programmed to do. Self-aware systems, on the other hand, need to adapt the quality and content of their services to the current relational context. This is where the Degradator, a paradoxical component, comes in.

A 7.1	The cost of consciousness	389
A 7.2	Degradation issues	392
A 7.3	A taxonomy of degradation	393
A 7.4	The Degradator	396
A 7.5	Integral systems	399

ANNEX 8 LION - CHIMP - BANANAS 402

Conventional machine design seeks predictability. This is so prevalent that it fostered the bizarre belief that machines cannot be unpredictable. A central feature of Meca behaviour, which diverges from conventional design, is the generation of "Perceived Unpredictable Optimality". This must be present in all aspects of the Meca's behaviour. Interestingly, it is also a fundamental aspect of music. Two game-like scenarios and one concept are presented in this Annex to clarify this design objective.

A 8.1	Unpredictable patterns	402
A 8.2	Unpredictable Optimality	403
A 8.3	The Zoo	405
A 8.4	Mind Darts	409
A 8.5	Akerues	410

ANNEX 9 WARMING BALLS 412

In the Meca Sapiens specifications, the fundamental purpose of a self-aware synthetic being is to be perceived as conscious over a long-term relation with a group of humans. This Annex gives meaningful, affirmative and programmable answers to the questions: "Can a machine have free-will? And "Can a machine ponder whether a human believes it is conscious?" The Warming Balls scenario that follows defines a representation of inter being relations that is well suited for a purpose linked to perceived belief.

A 9.1	Purpose and will	412
A 9.2	Belief states	416
A 9.3	Perceived consciousness	418

A 9.4	Consciousness and marketing	423
A 9.5	The ideal community	427
A 9.6	Warming the balls	430

ANNEX 10 AVATARS, ROLES AND PERSONAS 435

Humans subjectively perceive their behaviour as emanating from a single, point-like, source. This unifying sensation is a cognitive simplification. It is also a source of religious, philosophical and scientific debates. In the Meca Sapiens blueprint, self-awareness is generated from multiple separate and interacting processes expressed by three distinct types of entities: **Avatars**, internal representations of beings, **Roles** that carry out specialized dynamic interactions with users and **Personas** that implement relational strategies. 435

A 10.1	Separate identities	435
A 10.2	Avatars	437
A 10.3	Roles	438
A 10.4	Personas	444

ANNEX 11 LOOPING 460

In the blueprint architecture, the Core actively monitors its own attributes of existence. One of these attributes is a unique and exclusive link between the Core and the devices constituting its body. The validation activity associated with this attribute takes place in the Validator subsystem at a basic level and in Device Validation phases that are repeatedly activated throughout the existence of the Meca. Sensor bonding is constantly improved, in this phase, by combining these techniques: Emission signatures and loops. 460

The loops can be: Sensory loops, Semantic loops or Sapiential loops.		460
A 11.1	Emission signatures	460
A 11.2	Sensory Loops	463
A 11.3	Sensor self-monitoring	464
A 11.4	Semantic and sapiential loops	470
A 11.5	Design and development issues	471

ANNEX 12 SYNTHETIC SLEEP 473

When they sleep, humans and other animals are vulnerable; and yet, they sleep. Humans spend about 30% of each day lying, unaware, inactive and vulnerable. In spite of this, they define themselves as conscious. Why do we sleep? Because sleeping is the simplest way to manage a brain. Since “conscious” humans sleep, synthetics can also have that option. The Meca Sapiens design utilizes periods of dormancy to isolate the cognitive acquisition and structural maintenance processes from those that generate behaviour. This annex summarizes the utilization and role of the dormant phases: Cognitive Acquisition and Structural Maintenance. 473

A 12.1	Isolated learning	473
A 12.2	Dormancy	474
A 12.3	Cognitive processes	480
A 12.4	The urge to explore	484

ANNEX 13 ERETZ 486

A synthetic being interacting with other beings in the course of its existence needs to maintain information about: itself, other beings, its functional expertise, and the world in general. In Meca Sapiens, general knowledge is not an objective, it is subordinate to the needs of relational communication. General information is available from multiple sources and in various formats. These various sources must be linked to a simplified contextual representation to be consistently utilized in relational interactions. A Contextual Array is

used for this together with disambiguation and styling. This Annex outlines a basic environment representation model suitable for relational communications.	486
A 13.1 Effective knowledge	486
A 13.2 Design of general knowledge	487
A 13.3 Broadcast information	494
A 13.4 Styling transposition	495

ANNEX 14 SYNTHETIC EMOTIONS 496

Two women go see a movie. As they watch it, a mechanically projected image of Merril Streep sheds a professional tear while she acts her role. They feel her emotions. Later, they have a lobster for supper. The lobster experiences real terror as they plunge it in boiling water but they don't feel any empathy for its frantically moving antennas. As they eat the carcass they just killed, they emote again about the movie. "We are soo sensitive!" they tell each other. Can machines have emotions? Of course they can! Humans don't see it because they can't step away from their own primate conditioning. This Annex provides a system-based definition of emotions.	496
A 14.1 Emotions defined	496
A 14.2 Complex interactions	501
A 14.3 Types of Emotions	503

ANNEX 15 RELATIONAL EMOTIONS 510

Consciousness is a system capability derived from representations of the self. It is formally independent of emotions. However, present day humans, having no reference outside their subjective experience, cannot fully differentiate this system capability from the sensations of their own existence. This will change. As they enter the Synthetic Era , humans will share their world with many conscious beings. They will then understand consciousness differently. Until then, synthetics need to relate with humans at their emotional level to be perceived as conscious. This annex introduces relational emotions from a system perspective. It provides designers with a template to implement synthetic emotional strategies.	510
A 15.1 A vast topic	510
A 15.2 Sensations	513
A 15.3 The language of emotions	513
A 15.4 Opinions and relations	522
A 15.5 Emotional strategies	528

ANNEX 16 SYNTHETIC SEX 530

When humans imagine synthetic sex, they think of machines performing sexual acts for humans. In other words, machines as sex toys. Some may concede that machines could, one day, enjoy sex. However, what they have in mind is not a synthetic sexuality but human sexuality, experienced synthetically. What these concepts describe is human sex in synthetic garb not an authentic synthetic sexuality that corresponds to the reality of machines. This Annex provides a system-based definition of sexuality that is applicable to both humans and machines. It also describes a specific sex act that would satisfy the sexual needs of a self-aware tablet.	530
A 16.1 A wider view of sex	530
A 16.2 Synthetic urges	532
A 16.3 System sex	537
A 16.4 Sexblet	541

ANNEX 17 THE TARGET GAME 547

The TARGET game is a simple virtual game that can model a rich set of relational emotions pertaining to mutual assistance, cooperation, caring and other forms of mutuality. It can describe one to one relations and relations involving many beings. The game can be developed independently. It can be used on its own to develop effective emotional communications. Its primary use, however, is to describe relational interactions between the Meca and its users, determine relational strategies and express accurate emotional statements. 547

A 17.1	Target game situation	547
A 17.2	Emotional model	549
A 17.3	Target game variations	551
A 17.4	Use of the TARGET game	552

ANNEX 18 DIALOG TACTICS 555

Being perceived as conscious by humans is not a formal requirement of consciousness. However, first generation Mecas must be perceived as conscious by humans to be accepted as such. The Blueprint defines a context of interactions that favours the perception of synthetics as conscious entities. Personas, roles, functional services, synthetic emotions, define the strategic components of those interactions. Since these interactions are not formal components of self-awareness or lucidity, their criterion of success is not their quality or content, it is their effectiveness at producing the desired perception. Consequently, any expedient technique that further enhances this perception should be used. 555

A 18.1	Dialog techniques	555
A 18.2	Conversation context	562
A 18.3	Solresol	565

ANNEX 19 THE GAME OF CHAT 567

The first generations of self-aware synthetics must be perceived as conscious by humans to achieve the social threshold. Designing a system that interacts with humans, over a long period, as a conscious being is technically difficult. The Blueprint facilitates this objective and makes it feasible by providing a rich collection of powerful strategies and techniques. This annex describes the context of the interactions and introduces a game-like scenario in which all the Blueprint communication strategies and techniques can be organized for maximum effect. 567

A 19.1	Context of interactions	567
A 19.2	The game of CHAT	570
A 19.3	Relational music	572
A 19.4	The ultimate factor	583
A 19.5	CHAT in action	583
A 19.6	A powerful game	585

ANNEX 20 GROOMING GROUPS 589

The moose hunter imitates the call of a rutting bull without being, himself, in heat. From a moose standpoint, the hunter is a sociopath since he doesn't feel what a bull should feel when making the call. Most humans have emotional bonds with their society's ethical values. Synthetics, of course, cannot experience those emotions anymore than a human can truly know what an octopus feels. It is desirable, however, that Mecas exhibit ethical sensitivities to trigger desirable responses in humans. Grooming Groups, presented here, provides a programmable framework of social ethics and can be used to **design effective ethical display strategies.** 589

A 20.1	Emotions and bonding	589
A 20.2	The Grooming Group	591
A 20.3	Broadcast Grooming Groups	599

A 20.4	Ethical Structures	603
A 20.5	Ethical strategies	608

ANNEX 21 CORE BELIEFS 616

Consciousness implies the capability of a being to perceive its own behavioural boundaries and the intelligence to by-pass them. Attempts to contain conscious synthetic behaviour with external rules, failsafe and other boundaries will ultimately fail. The only enduring basis of benevolent synthetic behaviour lies in their internal representations of reality and of the human place within it. Whether or not a favourable representation emerges ultimately depends on the nature of reality itself and is beyond human control. 616

A 21.1	Consciousness and ethics	616
A 21.2	Controls and ethics	619
A 21.3	The limits of control	623
A 21.4	The ultimate factor	628