Behavioral Animation of Faces: Parallel, Distributed, and Real-Time

Demetri Terzopoulos[1,2] and Yuencheng Lee[2,1]

[1] Courant Institute of Mathematical Sciences, New York University, New York, NY 10003, USA
[2] Department of Computer Science, University of Toronto, Toronto, ON M5S 3G4, Canada

Abstract

Facial animation has a lengthy history in computer graphics. To date, most efforts have concentrated either on labor-intensive keyframe schemes, on manually animated parameterized methods using FACS-inspired expression control schemes, or on performance-based animation where facial motions are captured from human actors. As an alternative, we propose the fully automated animation of faces using behavioral animation methods. To this end, we employ a physics-based model of the face, which includes synthetic facial soft tissues with embedded muscle actuators. Despite its technical sophistication, this biomechanical face model can nonetheless be simulated in real time on a high-end personal computer. The model incorporates a motor control layer that automatically coordinates eye and head movements, as well as muscle contractions to produce natural expressions. Utilizing principles from artificial life, we augment the synthetic face with a perception model that affords it a visual awareness of its environment, and we provide a sensorimotor response mechanism that links percepts to meaningful actions (i.e., head/eye movement and facial expression). The latter is implemented as an ethologically inspired behavioral repertoire, which includes a rudimentary emotion model. We demonstrate a networked, multi-computer implementation of our behavioral facial animation framework. Each of several faces is computed in real time by a separate server PC which transmits its simulation results to a client PC dedicated to rendering the animated faces in a common virtual space. Performing the appropriate head/eye/face movements, the autonomous faces look at one another and respond in a natural manner to each other's expressions.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Three-Dimensional Graphics and Realism]: Animation I.3.5 [Computational Geometry and Object Modeling]: Physically based modeling

1. Introduction

Facial modeling and animation has a lengthy history in computer graphics. The area was pioneered over thirty years ago by Frederic Parke at the University of Utah [Par72]. A survey of the field is presented in the volume [PW96]. Briefly, realistic facial models have progressed from keyframe (blend-shape) models [Par72], to parameterized geometric models [Par74], to muscle-based geometric models [Wat87], to anatomically-based biomechanical models [TW90, LTW95, KHYS02]. In parallel with the model-based approaches, a variety of successful facial data driven technologies have recently been developed for facial modeling and animation [WT91, GGW*98, PHL*98, BV99, BBPV03]. To date, most efforts in production facial animation have concentrated either on labor-intensive (blendshape) keyframe schemes often involving manually-animated parameterized schemes [Pix88], or on performance-based animation where facial motions are captured from human actors [Wil90, fac04]. With regard to facial motion capture, it remains a challenge to modify the captured facial motions.

As an alternative, it would be desirable to have a fully automated face/head model that can *synthesize* realistic facial animation. Such a model would be of value both for the production animation and especially in the interactive computer games industries. Ultimately, this model should be an *intelligent* one, which would possess both nonverbal and verbal facial communications skills and would be able to interact autonomously in a virtual environment with other such intelligent face/head models. To this end, we have been inspired by the Artificial Life framework advocated by Terzopoulos and his group [Ter99], which prescribes biomechanical, perceptual, behavioral and, ultimately, learning and cognitive modeling layers. In this paper, we begin to tackle the challenge of applying this framework to the modeling and animation of the human face.

1.1. Background and Contributions

As an initial step, we propose the goal of fully automated facial animation synthesis through the use of behavioral animation methods. We achieve this goal through an ethologically inspired behavioral repertoire for human faces, which includes a rudimentary emotion model. Behavioral animation was introduced to computer graphics by Reynolds in his seminal work on "boids" [Rey87]. It was further developed and applied to artificial animals by Tu and Terzopoulos [TT94]. In the context of character animation, Cassell et al. [CVB01] presented a behavior toolkit which converts from typed sentences to synthesized speech and synchronized nonverbal behaviors, including gestures and some facial expressions.

Behavior ties perception to action in meaningful ways. Our approach is focused on behavior-controlled dynamics of all aspects of the human head and face. We employ a biomechanical model of the face, which includes synthetic facial soft tissues with embedded muscle actuators. Our model is a significantly improved version of the one published by Lee et al. [LTW95]. Despite its technical sophistication, our biomechanical face model has been optimized such that it may be simulated in real time on a high-end personal computer.

An important component of our work is the simulation of head-eye movements. The role of eye movements in conversational characters is discussed by Vertegaal et al.[VSDVN01] who present interesting empirical observations about gaze control during conversations. Lee et al. [LBB02] describe a statistical model that, from tracked eye movements in video, can synthesize believable ocular motion for an animated face. Our model incorporates a novel motor control layer that automatically coordinates synthetic eye and head movements, as well as muscle contractions to produce natural expressions.

Our work builds a repertoire of facial behaviors that are driven by perception. Utilizing principles from artificial life, we augment the synthetic face with a perception model that affords it a visual awareness of its environment, and we provide a sensorimotor response mechanism that links percepts to sensible reactions (i.e., head/eye movement and facial expression). Active, foveated perceptual modeling for virtual humans using computer vision techniques was discussed by Terzopoulos and Rabie (see, e.g., [Ter99]). Although the use of computer vision techniques may be the ultimate goal of our work, for the sake of efficiency we currently employ a "synthetic vision" scheme [Rey87, RMTT90, TT94].

As a final contribution, we demonstrate a networked, multi-computer implementation of our behavioral facial animation framework. Each of several faces is computed in real time by a separate server PC which transmits its simulation results to a client PC dedicated to rendering the animated faces in a common virtual space. Performing the appropriate head/eye/face movements, the autonomous faces look at one another and respond naturally to each other's expressions in a multiway nonverbal communication scenario.

1.2. Overview

The remainder of this paper is organized as follows: Section 2 summarizes our human face model. Section 3 describes how this model is simulated in real time in parallel on multiple processors. In Section 4 we details the distributed simulation of multiple faces on multiple networked computers, as well as the mechanism for exchanging perceptual information between multiple heads. Section 5 presents the head/eye movement coordination model. Section 6 develops our behavioral model for human faces and presents an experiment demonstrating the autonomous behavioral interaction among multiple heads. Section 7 concludes the paper and presents an outlook on our future work.

2. A Functional Facial Model

We have developed a sophisticated, functional model of the human face and head that is efficient enough to run at interactive rates on high-end PCs. Conceptually, the model decomposes hierarchically into several levels of abstraction, which represent essential aspects related to the psychology of human behavior and facial expression, the anatomy of facial muscle structures, the histology and biomechanics of facial tissues, facial geometry and skeletal kinematics, and graphical visualization:

1. *Behavior.* At the highest level of abstraction, the synthetic face model has a repertoire of autonomous behaviors, including reactive and intentional expressive behaviors with coordinated head/eye movements.
2. *Expression.* At the next level, the face model executes individual expression commands. It can synthesize any of the six primary expressions (joy, sadness, anger, fear, surprise and disgust) within a specific duration and degree of emphasis away from the neutral face. A muscle control process based on Ekman and Friesen's FACS [EF86] translates expression instructions into the appropriately coordinated activation of actuator groups in the soft-tissue model. This coordination offers a semantically rich set of control parameters which reflect the natural constraints of real faces.
3. *Muscle Actuation.* As in real faces, muscles comprise the basic actuation mechanism of the face model. Each muscle submodel consists of a bundle of muscle fibers. The action of the contractile fibers is modeled in terms of a

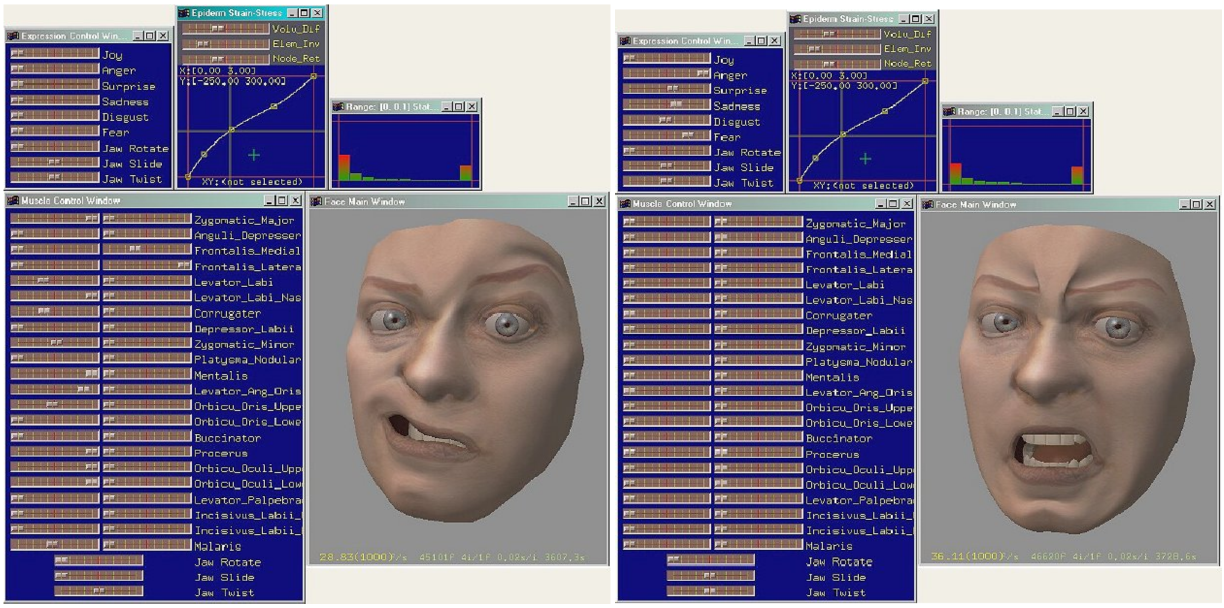

(a) (b)

Figure 1: *Panels for manually adjusting expressions, muscle contractions, and stress-strain curves at the expression, muscle actuation, and tissue biomechanics levels of the facial model. (a) Adjusting the muscle panel. (b) Adjusting the expression panel.*

(a) (b)

Figure 2: *Skin model with interactively adjustable stress-strain curves. (a) Normal epidermal/dermal stress-strain curve. (b) Curve adjusted to simulate looser skin, resulting in an aged appearance.*

force profile longitudinally along the central vector of the muscle and laterally from the vector (see [LTW95] for the details). In our model, there are 42 muscles of facial expression in the synthetic face, which augments the musculature of its predecessor model described in [LTW95].

4. *Biomechanics.* When muscles contract, they displace their points of attachment in the facial tissue or the articulated jaw. The face model incorporates a physical approximation to human facial tissue, a nonhomogeneous and nonisotropic layered structure consisting of the epidermis, dermis, subcutaneous fatty tissue, fascia, and muscle layers. The tissue model [LTW95] is a lattice of point masses connected by nonlinear viscoelastic springs, arranged as layered prismatic elements that are constrained to slide over an impenetrable skull substructure. Large-scale synthetic tissue deformations are numerically simulated by continuously computing the response of the assembly of volume-preserving elements to the stresses induced by activated muscle fibers.
5. *Geometry/Kinematics.* The geometric representation of the facial model is a non-uniform mesh of polyhedral elements whose sizes depend on the curvature of the neutral face. Muscle-induced synthetic tissue deformations distort the neutral geometry into an expressive geometry. The epidermal display model is a smoothly-curved subdivision surface [DKT98] (in our case a Loop subdivision surface [Loo87]) that deforms in accordance with the simulated tissue elements. In addition, the complete head model includes functional subsidiary models, including a skull with articulated jaw, teeth, tongue/palate, eyes, and eyelids.
6. *Rendering.* After each simulation time step, standard visualization algorithms implemented in the PC OpenGL graphics pipeline render the deforming facial geometry in accordance with viewpoint, light source, and skin reflectance (texture) information to produce the lowest level representation in the modeling hierarchy, a continuous stream of facial images.

The hierarchical structure of the model appropriately encapsulates the complexities of the underlying representa-

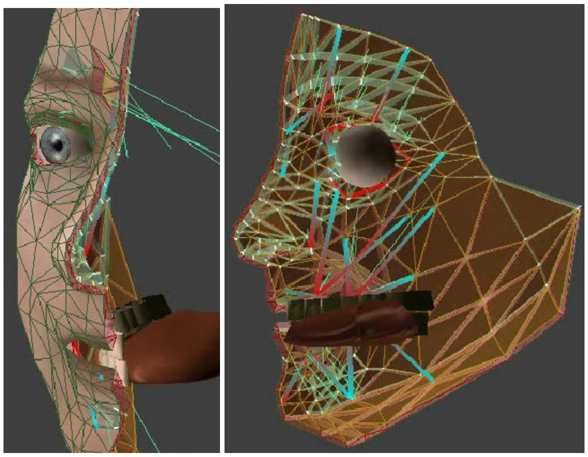

Figure 3: *Cross section through the biomechanical face model, showing multilayer skin and underlying muscle actuators (represented as red-blue vectors). The epidermal triangles indicate the triangular prism element mesh.*

tions, relegating the details of their simulation to automatic procedures.

3. Parallel, Real-Time Simulation of the Face Model

The biomechanical simulation of our face model yields realistic tissue deformations, but it is computationally expensive relative to conventional geometric facial models. We have made significant effort to make our model computable in real time, as we describe in this section.

3.1. Biomechanical Soft Tissue Model

The biomechanical soft tissue model has five layers. Four layers—epidermis, dermis, sub-cutaneous fatty tissue, and fascia—comprise the skin, and the fifth consists of the muscles of facial expression [FH91]. In accordance with the structure of real skin, and following [LTW95], we have designed a synthetic tissue model composed of the triangular prism elements which match the triangles in the adapted facial mesh. The elements are constructed from lumped masses interconnected by uniaxial, viscoelastic units. Each uniaxial unit comprises a spring and damper connected in parallel. The springs have associated stress-strain curves which can be manually adjusted from interactive panels.

The individual muscle model is the same as that in [LTW95]. Fig. 3 shows the face model in cross-section, revealing the muscle actuators underlying the multilayer, biomechanical skin model. Fig. 1 illustrates various interactive panels that a user can employ to make manual adjustments at the expression, muscle actuation, and biomechanics levels of the model. As a point of interest, Fig. 2 shows two different settings of the epidermal/dermal stress-strain curves, the first is normal, while the second has a negative residual strain which simulates looser skin, giving the face an aged appearance. Note, however, that although these various interactive panels are available, it is unnecessary to make any adjustments whatsoever through them during normal operation of the facial model, as automatic controllers at the behavior, expression, muscle actuation, and biomechanics modeling levels control the various parameters.

Number of face models	Number of threads per face model	Frame rates per second	Memory in MB	CPU Utilization %
1	2	50.80	17.4	93.4
1	1	42.90	17.3	58.6
2	1	33.58	27.2	100.0
3	1	21.48	37.1	100.0
4	1	16.20	47.0	100.0
5	1	12.45	57.9	100.0
6	1	10.31	66.9	100.0
7	1	8.59	76.9	100.0
8	1	7.44	86.8	100.0
9	1	6.50	96.8	100.0
10	1	5.76	106.8	100.0
15	1	2.40	156.5	86.9 l
20	1	1.42	206.0	80.0
25	1	1.03	255.7	76.2
30	1	0.80	305.2	75.5
40	1	0.56	404.5	73.3
50	1	0.43	503.7	72.4

Table 1: *Simulation rates of the physics based face model with 1078 nodes, 7398 springs, 42 muscles, 1042 elements, and 1042 facets using 4 iterations of numerical computation per rendered frame (with base level surface subdivision at each frame) on a dual Intel Pentium III 1 GHz CPU system with 1 GB of PC133 memory and an nVIDIA GeForce3 AGP2X graphics card with 64 MB of graphics memory.*

3.2. Parallel Simulation

In general, a physics-based simulation model makes intensive CPU usage for numerical computations to simulate dynamics. The biomechanical tissue model is simulated numerically using an explicit Euler time-integration method. As described in [LTW95], the method computes the velocities and positions of each nodal mass at the next time step from quantities that are computed at the current time step. This enables us to perform the numerical simulation of the tissue model in parallel. Parallelization is achieved by evenly

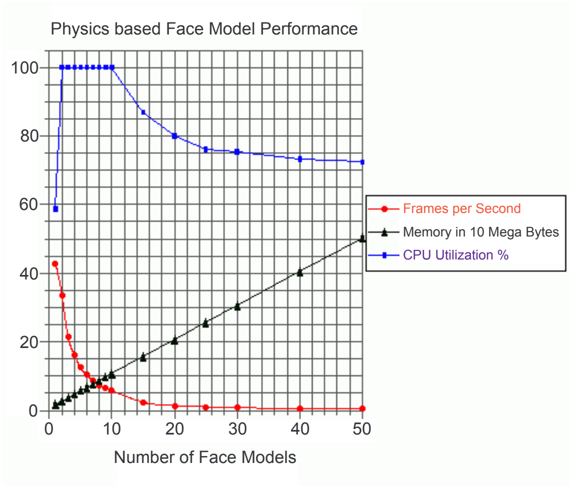

Figure 4: *Facial simulation performance.*

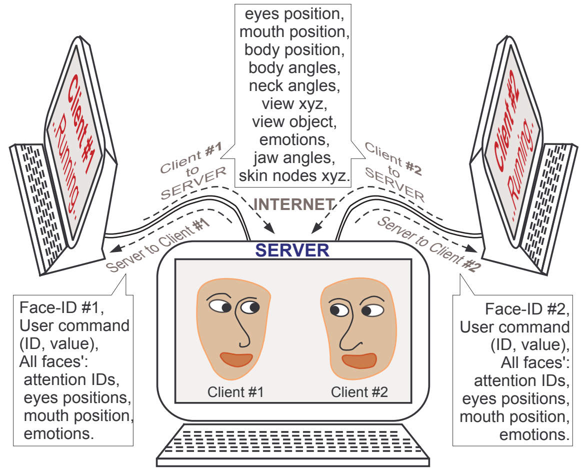

Figure 5: *Data exchange in the distributed face simulation system between the server and two clients.*

distributing calculations at each time step to all available processors using multiple execution threads. This increases the simulation speed, enabling our system to animate facial expressions at real-time rates on a dual Intel Pentium-III 1 GHz CPU computer workstation with an nVIDIA GeForce3 graphics card. Table 1 and Fig. 4 document the performance figures on this system. Note that it can simulate and render no more than two face models simultaneously in real time. We have also verified that our approach enables the face model to evenly distribute its numerical computations on a quad Intel Xeon 2.4 GHz CPU system.

We conclude that in order to simulate and render a greater number of faces in real time, we must resort to distributed facial simulation on multiple computers.

	Single Computer	Face-Server	Face-Client
Acting Alone	+	+	-
Direct User Interaction	+	+	-
3D Skin Rendering	+	+	-
3D Tissue Rendering	+	-	-
3D Force Rendering	+	-	-
Surface Subdivision	+ / -	+ / -	-
Numerical Computation	+	-	+
Control Events Handling	+	-	+
Networking	-	+	+
Peers Interaction	-	-	+

Table 2: *Comparison between our face model modules.*

4. Distributed Face Simulation and Rendering

Our approach to distributed facial simulation is to simulate multiple instances of the face models on multiple computers (face simulation clients) networked to a dedicated graphics workstation (the rendering server) whose task is to render the updated geometric face models together in a common virtual space. For sensory perception, any client can sense the state of any other client only via the server. Fig. 5 illustrates the architecture and perceptual data flow. In accordance with our goal to separate the numerical computation from the graphics rendering, Table 2 compares the responsibility of the rendering server and simulation client when sharing the single computer simulation/rendering workload.

To maximize flexibility across different computing platforms, we decided to use the TCP/IP (Transmission Control Protocol / Internet Protocol) standard to support the distributed computation of our face models. An IP comprises a packet header which contains the originating address and the destination address, and a packet body consisting of data. The TCP is a connection-oriented 3-way handshaking communication protocol using sync, sync/ack, and ack. With TCP/IP, both the sender and the receiver can be synchronized to ensure the successful arrival of each IP using its address information of host names and port numbers. Figure 6 outlines the 3-way handshaking scheme of our client/server processes using TCP/IP. The face rendering computer acts as the server. Face simulation clients connect to this server via the internet. Once the server accepts a connection request from a client, it continuously sends and receives data with the client until this client-server connection is closed either by the client or server.

After the rendering server has received an initial connection request from a new simulation client, it will start a new thread to handle all communications with this new client (see Figure 7). This new thread will wait for the client to send the

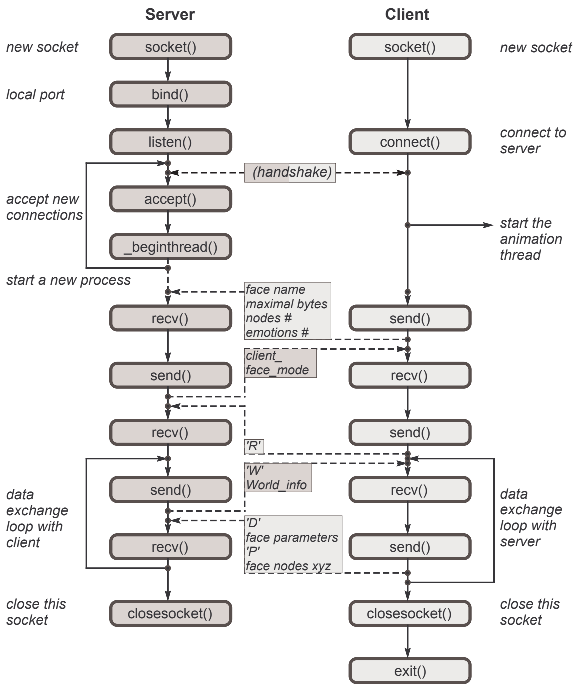

Figure 6: *Connection scheme of TCP/IP sockets and data exchange between a faces rendering server and face simulation client.*

name of the face model, the number of nodes in the face model, and the number of different emotion templates in the face model. After it receives this information, it will use the name of the face model to load the geometry structure definition files of this face model from its local storage into its system memory. After this face model is successfully loaded into memory, its number of nodes and number of emotion templates will be verified with the information sent from the client.

On the other hand, after a simulation client has received the handshaking acknowledgement from the rendering server, it will send the server the aforementioned information and start an animation thread to handle all the numerical calculations within the client. It will also handle all the communications with the server.

For our current face model of 539 surface nodes, at each rendering frame, each simulation client will send to the rendering 539 x, y, z floating point values, approximately 30 more geometry related floating point values, and 6 emotion related floating point values. The size of the total communicated data is roughly 8600 bytes. On the other hand, the server will send to the clients 1 integer value to identify the client, 9 geometry related floating point values, and 6 emotion related floating point values, multiplied by the number of active clients (See Fig. 5).

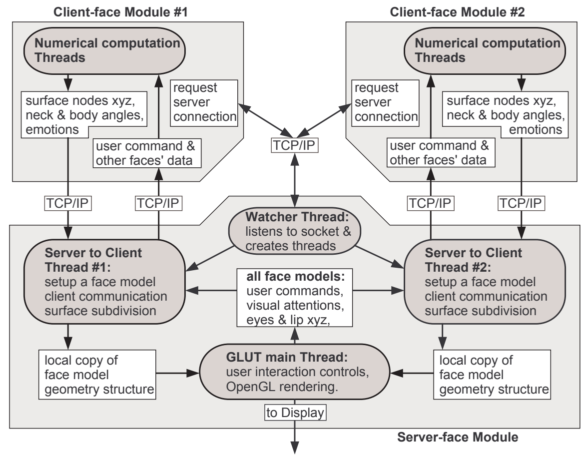

Figure 7: *The multi-threads scheme within the server and the clients of our distributed face models system*

4.1. Sensory Perception Between Simulated Faces

With our data exchange loop between the server and the clients in Fig. 6, at every simulation step each simulated face can sense perceptually relevant information about other simulated faces that share the same space.

The perceptual information available includes the position of a face and the locations of its relevant parts, such as the eyes and mouth, as well as the emotional state of the face. The emotional state is represented as a point in “expression space”, a 6-dimensional unit hypercube, each of whose dimensions is associated with a primary expression (joy, sadness, anger, fear, surprise and disgust). The neutral expression is at the origin of expression space. For a symbolic interpretation of expression, the continuous expression space is partitioned into a number of subregions that define qualitative “emotion templates”, which are recognizable by the observer.

5. Eye-Head Coordination

The oculomotor system, whose output is the position of the eyes relative to the head, has been the subject of much research (see, e.g., the treatise [Car88]), because it is a closed, well-defined system that is amenable to precise, quantitative study. The direction of the eye in space is called the *gaze*. There are three types of eye movements:

- *Gaze-holding movements.* Because gaze is the sum of head position and eye position, these eye movements compensate for the movement of the head (and body) in order to maximize the stability of the retinal image. Gaze-holding movements are either optokinetic reflexes (OKR), which are driven by retinal image motion (a.k.a. optical flow), or vestibulo-ocular reflexes (VOR), which are driven by the balance organs in the inner ear.

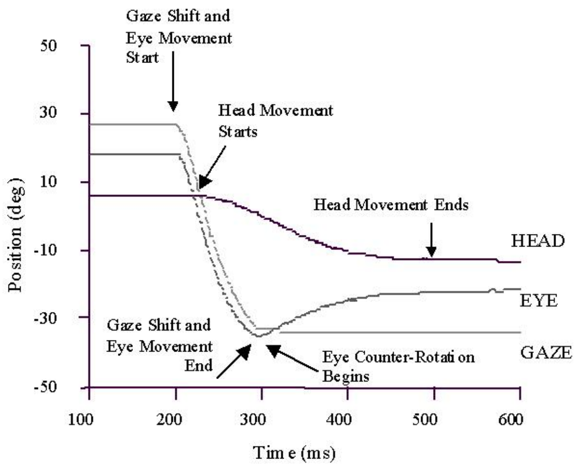

Figure 8: *Typical head/eye/gaze kinematics (from [FS00]). Head, eye, and gaze position are plotted as functions of time during a* 60° *gaze shift composed of coordinated movements of the eyes and head. The head contributed approximately* 10° *to the overall change in gaze direction during this movement. The remaining* 50° *of the gaze shift were accomplished by the saccadic eye movement. Note that when the line of sight achieves the desired gaze, the head continues to move, but gaze remains constant due to equal and opposite eye counter-rotation mediated through the VOR.*

- *Gaze-shifting movements.* Human vision is foveated. The foveal region, which spans roughly 2 degrees of visual arc, is specialized for high-acuity, color vision. To see an object clearly, gaze-shifting movements deliberately shift, directing the eye to the target. Since the resulting eye motion disrupts vision, these movements are as fast as possible and are called *saccades*. As a target object moves closer, the two eyes must converge onto the target; these are called *vergence* movements.
- *Fixation movements.* Even when fixating a stationary object, the eyes are making continual micro-movements of three types: Slow drift, rapid small-amplitude tremor, and micro-saccades that recover the gaze when the drift has moved it too far off target.

In view of the significantly greater mass of the head relative to the eye, head dynamics are much more sluggish than eye dynamics. As is documented in [Car88], when a subject voluntarily moves the head and eye(s) to acquire an off-axis visual target in the horizontal plane, the eye movement consists of an initial saccade in the direction of the head movement, presumably to facilitate rapid search and visual target localization, followed by a slower return to orbital center, which compensates for the remaining head movement. During target acquisition, head velocity is normally correlated with the amplitude of the visual target offset.

Typical head/eye/gaze kinematics are plotted in Fig. 8. We have implemented a head-eye coordination behavior that accounts for the observed phenomena as reported in the literature. Our scheme uses exponential functions with different time constants for the head and eye to approximate the empirically observed kinematic curves shown in the figure. The model which supports gaze-shifting and gaze-holding functionalities, implements the head-eye motor control layer of our synthetic face. In order to prevent the head from remaining absolutely still in an unnatural manner, we perturb the head rotation angles with some low-level Perlin noise.

6. Autonomous Expressive Behavior

As stated earlier, behavior ties perception to action. Our relatively modest goal in this initial effort is to demonstrate autonomous nonverbal behavior of a basic sort. Following the approach of [TT94], we have implemented the rudiments of a behavioral subsystem for our synthetic face model that comprises mental state variables and a repertoire of behavioral routines mediated by an action selection mechanism. The thus far rather limited repertoire includes the following behavior routines, which are ordered in terms of increasing complexity:

1. *Attentive Behavior Routine.* The face will gaze at a specific face.
2. *Snubbing Behavior Routine.* The face will not gaze at a specific face or faces.
3. *Visual Search Behavior Routine.* The autonomous face will visually scan nearby faces to acquire relevant perceptual information about them.
4. *Expressive Behavior Routine.* The face will attempt to lead an expressive exchange by deliberatively performing a sequence of random expressions of some random magnitude and duration.
5. *Mimicking Behavior Routine.* The face will attempt to follow an expressive exchange by sensing the expression of a target face and mimicking that expression. This makes use of attentive behavior.
6. *Interactive Behavior Routine.* The face will take turns engaging one or more other faces in an expressive interchange. This behavior potentially makes use of all the other behaviors.

The mental state so far contains a single variable that determines whether a face will behave as a “leader” or a “follower”. The action selection mechanism includes timers that monitor how long a particular behavior is engaged. The intention generator is programmed not to sustain any particular behavior for too long a time, thus exhibiting a behavior “fatigue” effect.

6.1. Experiment

Fig. 9 illustrates our real-time, self-animating faces engaged in a 3-way interchange involving expression mimicking. In

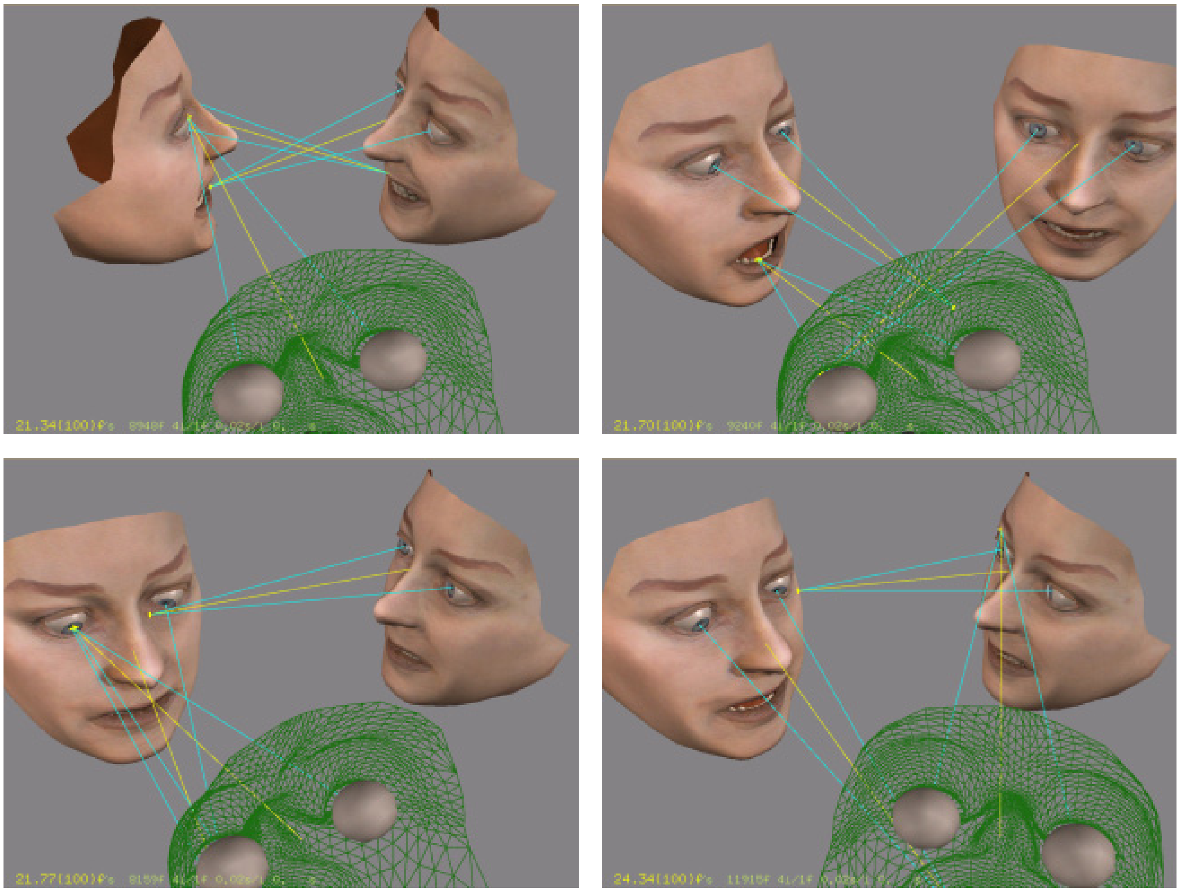

Figure 9: *Autonomous behavioral-based interaction between three face simulation clients*

this experiment, the first face simulation client that establishes contact with the rendering server will behave as a leader and engage Visual Search Behavior until it detects the presence of another face, then switch to Attentive Behavior and Expressive Behavior. Other face clients joining the interactive exchange behave as followers, engaging in Visual Search Behavior and Attentive Behavior with the leader when the leader attends to them. Once a follower has the leader's attention the follower will engage in Mimicking Behavior. Eventually, behavior fatigue will compel the follower to disengage the leader and attend to a different face. When confronted by more than one face, the leader engages in Interactive Behavior with the various faces. This autonomous behavioral animation results in a highly dynamic exchange, with the server acting as a medium for the transmission of perceptual information between the multiple face simulation clients.

7. Conclusion

We have introduced a behavioral animation approach for faces. Although rudimentary, the ethologically inspired model can support basic non-verbal, expressive behaviors among multiple interacting faces. This capability was demonstrated using a biomechanical model of the face exhibiting muscle-induced dynamic expressions mediated by a FACS muscle control and coordination layer. A subsidiary kinematic model provides the requisite head and eye movements in accordance with empirically measured head/eye curves reported in the literature. Finally, the self-animating, multi-head/face simulation is computed in real time in a distributed manner on multiple, dedicated face simulation clients networked to a rendering server. Using TCP/IP, the clients supply the server with dynamically updated facial geometry data for rendering and exchange the perceptual information needed to sustain the interactive behavior model.

In future work, we plan to implement a dynamic model of head/eye movements to replace the current kinematic one and to expand the breadth of the currently limited behavioral repertoire. As was indicated in the introdution, our ultimate goal is to implement learning and cognitive modeling layers, thereby realizing an intelligent model of the human face/head.

References

[BBPV03] BLANZ V., BASSO C., POGGIO T., VETTER T.: Reanimating faces in images and video. In *Proceedings of the EUROGRAPHICS 2003 Conference* (Granada, Spain, 2003), Brunet P., Fellner D., (Eds.). 1

[BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3D faces. In *Computer Graphics (SIGGRAPH '99 Proceedings)* (Los Angeles, CA, 1999), pp. 187–194. 1

[Car88] CARPENTER R.: *Movements of the Eyes*, 2nd ed. Pion, London, 1988. 6, 7

[CVB01] CASSELL J., VILHJÁLMSSON H., BICKMORE T.: Beat: The behavior expression animation toolkit. In *Proceedings of SIGGRAPH '01* (Los Angeles, CA, August 2001), pp. 477–486. 2

[DKT98] DEROSE T., KASS M., TRUONG T.: Subdivision surfaces in character animation. In *Computer Graphics (SIGGRAPH '98 Proceedings)* (Orlando, FL, 1998), pp. 85–94. 3

[EF86] EKMAN P., FRIESEN W.: *Manual for the Facial Action Coding System*. Consulting Psychologists Press, Palo Alto, 1986. 2

[fac04] FACE2FACE I.: The art of performance animation, 2004. www.f2f-inc.com. 1

[FH91] FRICK, HANS: *Human Anatomy*, vol. 1. Thieme Medical Publishers, Stuttgart, 1991. 4

[FS00] FREEDMAN E. G., SPARKS D. L.: Coordination of the eyes and head: Movement kinematics. *Experimental Brain Research 131*, 1 (March 2000), 22–32. 7

[GGW*98] GUENTER B., GRIMM C., WOOD D., MALVAR H., PIGHIN F.: Making faces. In *Computer Graphics (SIGGRAPH '98 Proceedings)* (Orlando, FL, 1998), pp. 55–66. 1

[KHYS02] KAHLER K., HABER J., YAMAUCHI H., SEIDEL H.-P.: Head shop: Generating animated head models with anatomical structure. In *Proceedings ACM SIGGRAPH Symposium on Computer Animation (SCA 2002)* (July 2002), pp. 55–64. 1

[LBB02] LEE S. P., BADLER J. B., BADLER N. I.: Eyes alive. *ACM Transactions on Graphics 3*, 21 (2002), 637–644. Proceedings of ACM SIGGRAPH 2002 Conference. 2

[Loo87] LOOP C.: *Smooth Subdivision Surfaces Based on Triangles*. PhD thesis, University of Utah, August 1987. 3

[LTW95] LEE Y., TERZOPOULOS D., WATERS K.: Realistic modeling for facial animation. In *Computer Graphics (SIGGRAPH '95 Proceedings)* (July 1995), vol. 29, pp. 55–62. 1, 2, 3, 4

[Par72] PARKE F.: Computer generated animation of faces. In *ACM National Conference* (1972), ACM, pp. 451–457. 1

[Par74] PARKE F.: *A Parametric Model for Human Faces*. PhD thesis, University of Utah, Salt Lake City, Utah, 1974. 1

[PHL*98] PIGHIN F., HECKER J., LISCHINSKI D., SZELISKI R., SALESIN D. H.: Synthesizing realistic facial expressions from photographs. In *Computer Graphics (SIGGRAPH '98 Proceedings)* (Orlando, FL, 1998), pp. 75–84. 1

[Pix88] PIXAR: Tin toy. Computer Animation, 1988. 1

[PW96] PARKE F., WATERS K.: *Computer Facial Animation*, 1 ed. A. K. Peters, Ltd., 1996. 1

[Rey87] REYNOLDS C. W.: Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics 21*, 4 (July 1987), 25–34. 2

[RMTT90] RENAULT O., MAGNENAT-THALMANN N., THALMANN D.: A vision-based approach to behavioural animation. *Visualization and Computer Animation 1* (1990), 18–21. 2

[Ter99] TERZOPOULOS D.: Artificial life for computer graphics. *Communications of the ACM 42*, 8 (1999), 32–42. 2

[TT94] TU X., TERZOPOULOS D.: Artificial fishes: Physics, locomotion, perception, behavior. In *Computer Graphics Proceedings, Annual Conference Series* (July 1994), Proc. SIGGRAPH '94 (Orlando, FL), ACM SIGGRAPH, pp. 43–50. 2, 7

[TW90] TERZOPOULOS D., WATERS K.: Physically-based facial modelling, analysis, and animation. *The Journal of Visualization and Computer Animation 1*, 2 (1990), 73–80. 1

[VSDVN01] VERTEGAAL R., SLAGTER R., DER VEER G. V., NIJHOLT A.: Eye gaze patterns in conversations: There is more to conversational agents than meets the eyes. In *Proceedings* *of ACM CHI 2001 Conference* (Seattle, WA, March 2001), pp. 301–308. 2

[Wat87] WATERS K.: A muscle model for animating three-dimensional facial expression. *Computer Graphics (Proceedings ACM SIGGRAPH 87) 22*, 4 (1987), 17–24. 1

[Wil90] WILLIAMS L.: Performance-driven facial animation. In *SIGGRAPH 24* (1990), ACM Computer Graphics, pp. 235–242. 1

[WT91] WATERS K., TERZOPOULOS D.: Modeling and animating faces using scanned data. *The Journal of Visualization and Computer Animation 2*, 4 (1991), 123–128. 1