

Methods for Online Management of AmI Capabilities Relative to Users' Goals

Ittai Flascher^{1,2}, Robert E. Shaw², Claire F. Michaels^{1,2}, and Oded M. Flascher^{2,3}

¹ Free University Amsterdam, Faculty of Human Movement Sciences, van der Bo-
echorststraat 9, 1081 BT Amsterdam, The Netherlands
{i.flascher, C_F_Michaels}@fbw.vu.nl
<http://www.fbw.vu.nl/index.htm>

² University of Connecticut, The Einstein Institute, 843 Bolton Road, Unit 1182, Storrs, CT
06269, U.S.A.
roberteshaw@aol.com
<http://www.ia.uconn.edu/einstein.html>

³ Visteon Corporation 17000 Rotunda Drive, Dearborn, MI, 48120, U.S.A.
oflasche@visteon.com
<http://www.visteon.com>

Abstract. Managing the various capabilities of computing environments to best support users' goals has proven a difficult problem in transportation systems and hand-held devices. In the case of mobile users, the goal of safe and efficient navigation is a persistent part of the users' context and therefore in the online decisions on what information and services to provide. We present a feasibility test of general methods for measuring and predicting actors' goal-directed performance, and outline their use in effecting decisions with regard to initiating and halting interactions with users, anticipation of users' needs, and the evaluation of Ambient Intelligence designs.

1 Introduction

A current challenge to the development of the envisioned seamless integration of users' computational and ecological (physical) environments is to manage all the available capabilities of Ambient Intelligence (AmI) to meet the demanding physical goals. In cars, aircrafts, and hand-held computing devices the need to manage the flux of available information and the use of communications and infotainment during operations has been most pressing [4],[5]. Several information management systems have been recently developed, and some related products are expected to reach the market in a few years [9],[17].

Underlying problems such as deciding when to initiate and halt interaction with the user, anticipating users' actions, and evaluating the ability of designs to support goal-directed actions still persist however, and are the current subject of intense research and development efforts [2].

Further challenges arise in the need to achieve AmI for mobile users across tasks and domains (e.g., home, workplace, a car). Current methods of information management are therefore required to become increasingly general to handle a spectrum of goals that may arise in users' everyday lives, as well as be adaptable to each individual's changing needs¹.

In the following we briefly outline some current approaches to information management and identify where the proposed methods can be of service in overcoming the above challenges.

1.1 Current Challenges in AmI Information Management

Initiation and Cessation of AmI Activities. Many currently available approaches to online information management effect a decision on questions such as whether to pass an incoming phone call to the user while the user is walking in a busy street in the following manner. First, the automated manager assembles all the available performance data (e.g., speed of vehicles, the average reaction-time of the user, workload etc.). It then identifies the most relevant measures for the achievement of the goal (e.g., reach the office quickly and safely). Finally, it decides on how to combine those into an overall measure of performance by which the current "state" of goal achievement is determined [e.g., 8]. Such a measure forms the basis for AmI decisions on initiating, halting, and prioritizing interactions with the user.

Partial measures, however, such as task duration, number of mistakes, among many others cover different aspects (dimensions) of performance and require a method to reconcile time, cognitive load, force, and number as components of goal achievement. Methods that can robustly quantify the connection between users' dynamics and constraints of different goals are needed.

Evaluation of AmI Designs. Currently available approaches attempt to establish the value of a given design through both subjective and cognitive measures (e.g., "expert" evaluation, users questionnaires), as well as by using partial performance measures. The proliferation of measures, however, opens door for a given design to be shown as better than a competitor by some of the measures and worse by others. Therefore, an overall measure of goal-relevant performance is additionally required by designers of AmI.

Anticipation of Goal-Relevant Performance. Prediction of users' performance on partial performance measures can be of great value. For example, predicting the effect that a phone ring will have on a given walker's deviation from a straight line can

¹ Much research is currently dedicated to developing interfaces that allow users to explicitly specify their goals to the AmI environment more easily [e.g., 15]. Other efforts are dedicated to developing automated AmI systems that can infer users' goals from their gestures and expressions [e.g., 5]. In this presentation we assume a goal has been specified and tackle the problem of managing AmI capabilities to support users in its achievement.

serve in the decision of whether to signal the user while the user is about to cross a street. Goals, such as reaching the office as quickly as possible, however, require a more general prediction. Predictions of whole dynamical paths are necessary, for example, to reschedule an incoming call to a later time or situation that would be optimal from the perspective of achieving the user's goal.

In the following we outline the proposed methods and measures, demonstrate their potential utility in solving the discussed problems in Aml implementation, and finally, we present empirical results from a recent experiment testing the feasibility of the framework.

2 Methods of Intentional Dynamics

2.1 Rationale

In the approach we take, goals are treated as a set of constraints on the outcome and/or process of performing tasks. For example, a person in the bedroom intending to go to the kitchen faces the task of walking (i.e., transportation). Getting to the kitchen as quickly as possible (i.e., minimal duration) sets a constraint on the task of walking and the observed dynamics will reflect that change to some degree. If a method can be found by which to quantify the influence (coupling) goal-constraints have on the dynamics, it will open three possibilities: First, a measure sensitive to every constraint on the dynamics should have properties by which determine when the actor is meeting the goal in all respects. Second, by the same argument the measure might also open the way to rank goal-performance for complete processes. Third, such an encompassing performance measure may reveal predictable regularities in goal-constrained dynamics where partial ones could not.

2.2 Methods²

To make the calculations explicit, we present an experiment where four volunteers controlled the motion of a graphically displayed sphere with a force-feedback joystick (i.e., a forcestick) [11].

The goal of the task (conveyed verbally) was to bring the controlled sphere to coincide with the target sphere in the shortest duration they can manage on every trial. Participants pressed a trigger to "release" the sphere and terminate the trial. Each participant performed ten sessions (in separate meetings) each consisting of five hundred trials. Data from a trial is termed a path. It consists of a sequence of samples of the forcestick's handle positions along two linear dimensions in time,

² In the following we give but a brief and intuitive outline of the methods. The reader is referred to [10] for further details on the approach. For further details on Simulated Annealing and related Monte Carlo techniques see e.g., [1,16].

$[q] = \{q_0, \dots, q_c, \dots, q_f\}$, where $0 \leq c \leq f$ stands for the current time-slice and f the final one in a completed path.




Fig. 1. The experimental setup. The graphics display (*left*), and Immersion's Impulse Engine 2000 force-feedback joystick (*right*)

For the purpose of presentation, the task we chose is simple relative to activities found in practice. Nevertheless, the task does involve all the basic components of the general transportation problem and the methods we present are applicable to the more complex cases.

Initiation and Cessation of AmI Activities. To support users' goal achievement, AmI requires a measure by which to determine online whether users' actions do not fulfill all the goal constraints and therefore assistance should be initiated. The same measure should also indicate when performance is adequate so that assistance could be properly terminated. In the following we outline the two computational stages by which the required indicator is achieved.

1. Path identification

The first step in the determination of whether a participant's performance satisfies the goal of the task is to identify at each time-slice (sample) all the paths that can reach the target and satisfy all the constraints. To achieve that we use a *Simulated Annealing* algorithm, which is a general method of sampling high-dimensional spaces [12]. The method is particularly advantageous in identifying paths when the number of constraints is large, and can do so for constraints originating from essentially any task.

The algorithm starts its search for goal-paths by constructing a random path. It then constructs a new path from the old one by randomly changing (perturbing) the old positions on some of the time-slices. This first out of two stages of the algorithm is termed the *generating* stage. Any known user or environmental constraints on path creation are treated here. For example, if the average maximal speed for the actor is known through any means, the algorithm directs its search to regions of the path space that include realistic paths for the actors' current capabilities.

In the *acceptance* stage, the algorithm selects which one of the paths will be recorded in a frequency table of goal-paths according to the following rule: if the new path is of shorter duration than the old, as the goal required, then the new path's counter is increased by one. If the new path is of longer duration it is not always rejected, rather, it is put to the following test:

$$\text{If, } \exp\left(\frac{-t_{old} + t_{new}}{\sigma_t}\right) > \text{rand}[0,1), \text{ then accept.}$$

The test is known as the *Metropolis criterion* [14], and in the case of our goal says: if the new path is not "much longer" in time than the old one then it will still be accepted. "Much longer" is quantified in the denominator by the standard deviation of observed trial duration for that actor, and the relative probability of acceptance is given by the negative exponent of the ratio. If the new path is still rejected the old path is recorded in the frequency table.

When set up properly, the algorithm approximates to a high degree the globally minimal (goal) path distribution within a few thousand samples. Correspondingly, the frequency table is divided by the number of samples yielding the probability distributions of goal-paths scaled to the known constraints on the user and the environment.

This stage of the analysis provides a set of paths in space and time that satisfy the goal to different degrees and may be executed by the user. In step 2. we develop a measure that indicates which path the user is actually following and ranks the feasible alternatives that Aml may promote.

2. Dynamics quantification

As all the goal and task constraints are implicit in the simulated paths, we need to quantify the dynamics that will be involved in producing such goal paths from the current time-slice. To get an intuition into what "quantifying dynamics" means in this context we first recall that the goal was to reach the target in minimal duration. We would like to consider all the forces and energy that might lead to the violation of the goal-constraints. For example, if the sphere controlled by the actor does not travel along the shortest distance to the target it will not be reached in minimal time. Similarly, if the sphere changes directions or speed unnecessarily, travel time will be prolonged. The same applies to any other goal (e.g., transport with minimal accelerations). We need to quantify all those relevant dimensions into a single scalar.

To achieve that, we sum three *action* terms along the mean simulated path $\left[\mathbf{q}^* \right] = \left\{ \mathbf{q}_c, \dots, \mathbf{q}_f^* \right\}$, from the current time-slice to the final time-slice when the sphere would have reached the target³,

$$S^*[c] = \sum_{t=c+1}^{t=f^*} \left[\frac{m}{2} \cdot \|\dot{\mathbf{q}}_t\|^2 \right] \cdot \Delta t + \sum_{t=c+1}^{t=f^*} \left[F_t^\perp \cdot \|\mathbf{q}_t\| \right] \cdot \Delta t + \sum_{t=c+1}^{t=f^*} \left[F_t^\parallel \cdot \|\mathbf{q}_t\| \right] \cdot \Delta t. \quad (1)$$

Action is the highest level dynamical variable and can be computed in several ways. The first term on the right side of the equation computes the action associated with the kinetic energy of the sphere's motion (where $\|\dot{\mathbf{q}}_t\|$ is the speed). The second term computes the action associated with the component of the resultant force signifying changes to the direction of motion (where $F_t^\perp$ is the Normal component of the force, and $\|\mathbf{q}_t\|$ is the distance along the path). The third quantifies the action arising from the changes to speed along the path. We term the sum of these components *prospective* action.

Similarly, at each time-slice we compute the amount of action already exerted up to the current one along the path traveled:

$$S[c] = \sum_{t=0}^{t=c} \left[\frac{m}{2} \cdot \|\dot{\mathbf{q}}_t\|^2 \right] \cdot \Delta t + \sum_{t=0}^{t=c} \left[F_t^\perp \cdot \|\mathbf{q}_t\| \right] \cdot \Delta t + \sum_{t=0}^{t=c} \left[F_t^\parallel \cdot \|\mathbf{q}_t\| \right] \cdot \Delta t. \quad (2)$$

We finally arrive at the quantity we were after by summing the *prospective* and *retrospective* actions into *generalized action* (GA) at time-slice c :

$$\tilde{S}[c] = S^*[c] + S[c]. \quad (3)$$

It is the constructed property of this quantity that serves as an indicator of the state of performance relative to the goal. More explicitly, in constructing the simulated paths we have essentially set up a new tracking task at each time-slice. Therefore, as long as the tracking is precise from one time-slice to the next the retrospective and prospective actions are complements and leave the values of GA invariant. In other words, if the actions of the user do not bring about the necessary dynamics required to satisfy the goal, invariance is not maintained. Such an indicator can therefore be used by AmI to determine when interventions are needed and when they are no longer required.

³ In general, there is a fourth term associated with a change in the orientation of the sphere, (i.e., a rotation around an axis). In our setup, the forcestick does not allow that degree of freedom.

For presentation purposes we set the target at $x = y = 40\text{mm}$, and clamp the maximal speed $\dot{q}_{\max}$ of the forcestick at 4mm/hundredth of a second along each axis (i.e., $\dot{q}_{\max} \approx 5.66$ in the Euclidean sense). By further setting the mass of the simulated object at 2 units, we get GA constant at 416.0.

Table 1. Action values for each time-slice in a goal-path

Time-slice	x,y	GA	Retro	Prosp
0	0,0	416.0	0.00	416.0
1	4,4	416.0	128.0	288.0
2	8,8	416.0	160.0	256.0
3	12,12	416.0	192.0	224.0
4	16,16	416.0	224.0	192.0
5	20,20	416.0	256.0	160.0
6	24,24	416.0	288.0	128.0
7	28,28	416.0	320.0	96.0
8	32,32	416.0	352.0	64.0
9	36,36	416.0	384.0	32.0
10	40,40	416.0	416.0	0.0
TGA		4,576		

Table 2. Constraint violation and the breakdown of invariance

Time-slice	x,y	GA	Retro	Prosp
0	0,0	416.0	0.0	416.0
1	4,4	416.0	128.0	288.0
2	8,8	416.0	160.0	256.0
3	12,12	416.0	192.0	224.0
4	16,16	416.0	224.0	192.0
5	20,20	416.0	256.0	160.0
6	24,24	416.0	288.0	128.0
7	28,28	416.0	320.0	96.0
8	32,32	416.0	352.0	64.0
9	35,35	400.0	364.0	36.0
10	39,39	400.0	404.0	4.0
11	40,40	400.0	400.0	0.0
TGA		4,944		

Invariance disappears for any violation such as the slowdown occurring on time-slice 9 in Table 2 below. Invariance returns as all the constraints are met and the motion is at maximal speed in the direction of the target from the tenth time-slice.

As we can see, this invariance of GA under goal-directed dynamics gives the sought after indicator to determine whether participants satisfy their goal while in the process.

Evaluation of AmI Designs. The sum of GA, Total-Generalized-Action (TGA) for a complete path j ,

$$\widehat{S}_j = \sum_{c=0}^{c=f^*} \widetilde{S}[c], \tag{4}$$

is minimal for goal-paths as can be seen from the previous tables and the next one.

Table 3. TGA calculations⁴

Time-slice	x,y	GA	Retro	Prosp
0	0,0	416.00	0.00	416.00
1	4,3	398.66	100.00	298.66
2	6,5	386.85	104.13	282.73
3	10,7	421.46	144.29	277.17
4	12,11	438.67	189.59	249.08
5	16,15	438.67	248.80	189.87
6	18,18	443.74	259.52	184.22
7	21,22	445.79	294.27	151.52
8	25,24	482.30	330.29	152.01
9	29,27	470.77	366.91	103.86
10	33,31	470.77	411.70	59.07
11	37,35	470.77	443.70	27.07
$f^* - 1$	40,39	470.77	472.50	1.74
f^*	40,40	470.77	470.77	0.00
TGA		6,226		

Therefore, TGA is an overall measure of success in satisfying the goal for a complete process. As the constraints of the goal are violated TGA increases and can therefore serve to rank the observed paths according to their merit in goal achievement. Comparing the effectiveness of competing AmI designs in facilitating users goal-relevant behaviour can therefore be carried out through the TGA measure.

⁴ In case the sphere is triggered away from the target the simulation completes the path; hence f^* instead of f .

Anticipation of Goal-Relevant Performance. Predictions of goal-directed behaviour may arise from knowledge of human cognitive and physical abilities and constraints. In the following we present a test of feasibility for a complementary approach to the human-centered approach. In the goal-centered approach we search for a principle capturing the predictable regularities in actors' task dynamics under goal-constraints. More specifically, we would like to formulate a principle that predicts for any given actor the probability distribution of TGA (i.e., the relative frequency of occurrence in repeated experiments). We formulate a "least-TGA principle" given by the Boltzmann distribution:

$$\Pr\{\hat{S}_j = \hat{s}\} \equiv p_j = \frac{1}{Z(\omega)} \exp\left(-\frac{\hat{S}_j}{\omega}\right), \quad (5)$$

where,

$$Z(\omega) = \sum_j \exp\left(-\frac{\hat{S}_j}{\omega}\right), \quad (6)$$

and,

$$\omega \equiv \sigma_s \cdot T. \quad (7)$$

In words, the probability of observing a path (in repeated experiments) is an exponential function of the (negative) value of TGA (i.e., $\hat{s}$); the higher the value of TGA for a path, the less likely the path is.

In addition to TGA, there are two quantities in the denominator affecting the distribution of TGA.

$$\sigma_s = \sqrt{\frac{1}{N} \sum_{j=1}^N \left(S_j[f] - \langle S_j[f] \rangle \right)^2}, \text{ where } \langle \rangle \text{ signifies the mean (average), is the stan-}$$

dard deviation of (retrospective) action distribution observed for a participant up to the time of prediction (N trials)⁵.

The standard deviation of the observed action distribution serves as the standard unit by which we measure TGA, and $0 < T < \infty$ is a multiplying factor which estimates the magnitude of the effect the goal-constraints had on the actor's performance (i.e., goal-coupling strength). More explicitly, when $T \rightarrow \infty$, the effect of the minimization principle disappears and the resulting distribution is the consequence of a random-walk under the constraints of the generating stage of the algorithm alone. As $T \rightarrow 0$, the shape of the distribution shifts towards the exponential of the Boltzmann distribution as can be seen by comparing the next two figures. Improving skill quantified by the generating parameters (e.g., maximal observed speed) may allow a user to produce paths with smaller TGA values. The observed change in the distribution due

⁵ In the implementation presented we measure the action after each trial (i.e., at the final time-slice f).

to that is a displacement towards the origin. T on the other hand, influences the relative frequencies with which paths with smaller values are performed. It is that type of improvement that is associated with developing goal-relevant expertise and gauged by T .

An observed TGA distribution of an experimental session (500 trials) is plotted in red. Given the large value of T , the predicted distribution in blue gives the random-walk result.




Fig. 2. $T \cong 5,000$

At lower values, the predicted distribution changes its shape towards the exponential.





Fig. 3. $T \cong 500$

T is therefore estimated by a separate Simulated Annealing algorithm searching for a value which brings the *means* of the distributions to coincide.

The feasibility test of the above principle should boost confidence in the approach by showing that the least-TGA principle both predicts the data of every participant to a large extent, as well as by demonstrating that participants converge to the Boltzmann equilibrium distribution of that measure as the number of session increases.

1. *First experimental hypothesis, $H_1^{(1)}$: $\Delta \gg 0$.*

Using only three generating parameters of maximal speed observed, mean reaction-time to trial onset, and the standard deviation of final (Euclidean) distance from the target, close to 60% of the total distribution of every session of every participant in the study were predicted. Due to space constraints we show only a sequence of a few sessions of one participant in the following figures.






Fig. 4. A sequence of predicted sessions

We use the percentage similarity measure, also known as *Weitzman's Δ* (delta) measure [6], to quantify the level of prediction. As its name suggests, the measure yields the percentage of the observed distribution intersected (overlapped) by the predicted one,

$$\Delta = \sum_i \min(p_i^{obs}, p_i^{pred}), \quad (8)$$

where i are the values of TGA for which both the observed and predicted distributions have probability larger than zero.

As a first approximation, the approach seems promising, predicting most of every session in the study. Much improvement can be expected with the introduction of more elaborate path construction schemes and additional constraining parameters.

2. *Second experimental hypothesis*, $H_1^{(2)}: \underset{N \rightarrow \infty}{D} \rightarrow 0$.

When researchers use the simulated annealing algorithm as an integration or simulation tool they try to make sure the algorithm converges to the globally optimal (equilibrium) distribution they seek. One of the components in achieving that result is to construct paths at the generating stage which are independent of one another. The (Metropolis) acceptance criterion can then make sure that the path sequence (chain) converges to the optimal distribution. In our use of the algorithm as the model for users' *learning* process, the paths are strongly dependent and convergence is not guaranteed. Participants in our case are in charge of constraining their control to meet the goal requirements as much as they can, given their *skill*-level measured by the generating parameters (e.g., maximal observed speed).

The second experimental hypothesis is therefore that participants will show a decreasing distance from the predicted equilibrium distribution (i.e., the optimal distribution) as the number of trials (sessions) is increased. In the next figure we compare the red observed distribution with the black equilibrium distribution generated (sampled) at the same level of standard deviation $\sigma_{\bar{s}}$ (i.e., bin-size along the horizontal axis). The latter distribution signifies the best goal-relevant performance this participant may achieve given his/her current sensitivity to the goal-constraints.




Fig. 5. The observed and the equilibrium distributions

To test whether participants' distributions show convergence to the Boltzmann equilibrium distribution specified by our principle, we measure for each of them the (variation) distance of the observed distribution from the $\pi(\sigma_{\bar{s}})$ distribution at each session [e.g., 7]:

$$D = \frac{1}{2} \sum_i \left| \pi_i(\sigma_{\bar{s}}) - p_i^{obs} \right|. \quad (9)$$

Distance values between the distributions were fitted with a least-square line. The slope of that line was tested through a reshuffling technique [13] to statistically de-

termine whether it is significantly different from zero. As can be seen in the figures below, convergence of modest rate was detected for three out of the four participants.




Fig. 6. Convergence assessment

These preliminary results are quite promising in showing good data fits and the possibility of convergence. Of course, further work is required and testing the methods' performance on more realistically complex tasks is necessary. However, the evidence is quite strong that the approach is feasible and that further work is warranted.

3 Summary of Contributions

Methods were outlined for the solution of three general problems currently impeding the development of management systems of AmI. We outlined the methods' execution of decisions on the initiation and cessation of AmI interventions, overall design assessment through the evaluation of the level of goal achievement by the TGA measure, and the prediction of goal-directed performance. The methods and measures demonstrated are generally applicable and are particularly suited for adapting to different users' changing capabilities and goals. If shown valid for the complex tasks found in practice, the offered methods promise to facilitate the solution to some of most hindering problems in AmI's development.

References

1. Aarts, E. H. L., and Korst, J. H. M. Simulated Annealing and Boltzmann machines. John Wiley & Sons, New York. (1989)
2. Brookhuis, K. A., de Waard, D., and Fairclough, S. H. Criteria for driver impairment. *Ergonomics*, 46, 5 (2003), 433–445
3. Camurri, A., Lagerlöf, I., and Volpe, G. Recognizing emotion from dance movement: comparison of spectator recognition and automated techniques. *Int. J. Human-Computer Studies*, 59, (2003), 213–225
4. Chávez, E., Ide, R., and Kirste, T. Interactive applications of personal situation-aware assistants. *Computers & Graphics*, 23, (1999), 903–915
5. Chen, F-S., Fu, C-M., and Huang, C-L. Hand gesture recognition using a real-time tracking method and hidden Markov models. *Image and Vision Computing*, (In press)
6. Clemons, T. E., & Bradley Jr., E. L. A nonparametric measure of the overlapping coefficient. *Computational Statistics & Data Analysis*, 34 (2000). 51–61
7. Denuit, M., and Bellegem, S. van. On the stop-loss and total variation distances between random sums. *Statistics & Probability Letters*, 53 (2001). 153–165
8. de Waard, D., Hernández-Gress, N. and Brookhuis, K. A. The feasibility of detecting phone-use related driver distraction. *Int. J. Vehicle Design*, 26, 1 (2001). 85–95
9. EVAID. Trademark of QinetiQ Limited, Farnborough, Hants, GU14 0LX U.K. (2003). <http://webdb4.patent.gov.uk/tm/number?detailsrequested=C&trademark=2299745>
10. Flascher, I. Goal-centered approach to the measurement of human-systems performance. Ph.D. dissertation. University of Connecticut. (In press)
11. Impulse Engine 2000, Software Development Kit, release 4.2, (January, 2001). Immersion Corp.
12. Laarhoven, P. J. M. van, Aarts, E. H. L. Simulated Annealing: Theory and applications. Reidel, Dordrecht. (1987)
13. Lunneborg, C. E. Data analysis by resampling: Concepts and applications. Duxbury Press, USA. (2000)
14. Metropolis, M., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. Equations of state calculations by fast computing machine. *Journal of Chemical Physics*, 21 (1953), 1087–1092
15. Prekop, P., and Burnett, M. Activities, context and ubiquitous computing. *Computer Communications* 26, (2003), 1168–1176.
16. Robert, C. P., and Casella, G. Monte Carlo statistical methods. Springer-Verlag, New York. (1999)
17. Wood, C., Leivian, R., Massey, N., Bieker, J., and Summers, J. Driver Advocate™ Tool. In *Proceedings of Driving Assessment 2001: International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*, Aspen, Colorado, August 14–17, 2001