Optimization Strategies for Cognition and Autonomy in Mixed Human-Robot Teams

Vaibhav Srivastava and Francesco Bullo

Center for Control,
Dynamical Systems & Computation
University of California at Santa Barbara
http://motion.me.ucsb.edu

AFOSR-DSI Data-to-Decisions & Autonomy Workshop
RMIT University, Melbourne, Australia, July 9th and 10th, 2012

Big Picture: Human-robot decision dynamics

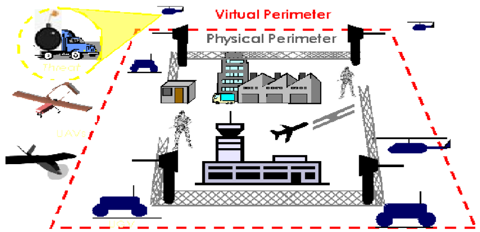

Uncertain environment surveyed by human-UAV team (Courtesy: Prof. Kristi Morgansen)

UAV surveillance (Courtesy: http://www.modsim.org/)

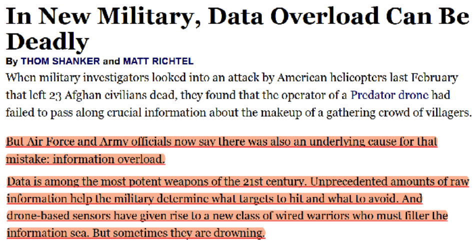

In New Military, Data Overload Can Be Deadly

By THOM SHANKER and MATT RICHTEL

When military investigators looked into an attack by American helicopters last February that left 23 Afghan civilians dead, they found that the operator of a Predator drone had failed to pass along crucial information about the makeup of a gathering crowd of villagers.

But Air Force and Army officials now say there was also an underlying cause for that mistake: information overload.

Data is among the most potent weapons of the 21st century. Unprecedented amounts of raw information help the military determine what targets to hit and what to avoid. And drone-based sensors have given rise to a new class of wired warriors who must filter the information sea. But sometimes they are drowning.

http://www.nytimes.com/2011/01/17/technology/17brain.html

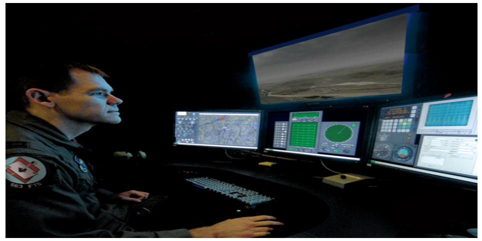

A surveillance operator (Courtesy: http://www.modsim.org/)

Two Critical Issues

Photo courtesy: The Wall Street Journal

Optimal information aggregation

- Which source to observe?
- Efficient search and detection
- Routing for evidence collection

Optimal information processing

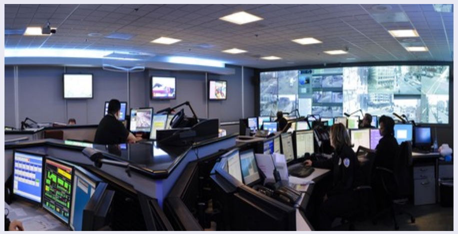

- Optimal time allocation?
- Optimal streaming rate?
- Optimal number of operators?

Cognition & Autonomy Management System (CAMS)
to optimize human-robot team objective

Incomplete Literature Review

Human Decision Making

R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen. The physics of optimal decision making: A formal analysis of performance in two-alternative forced choice tasks. *Psychological Review*, 113(4):700–765, 2006

R. W. Pew. The speed-accuracy operating characteristic. *Acta Psychologica*, 30:16–26, 1969

Control of Queues

O. Hernández-Lerma and S. I. Marcus. Adaptive control of service in queueing systems. *Systems & Control Letters*, 3(5):283–289, 1983

S. Ağrali and J. Geunes. Solving knapsack problems with S-curve return functions. *European Journal of Operational Research*, 193(2):605–615, 2009

Queues with human operator

K. Savla and E. Frazzoli. A dynamical queue approach to intelligent task management for human operators. *Proceedings of the IEEE*, 100(3):672–686, 2012

L. F. Bertuccelli, N. Pellegrino, and M. L. Cummings. Choice modeling of relook tasks for UAV search missions. In *American Control Conference*, pages 2410–2415, Baltimore, MD, USA, June 2010

N. D. Powel and K. A. Morgansen. Multiserver queueing for supervisory control of autonomous vehicles. In *American Control Conference*, Montréal, Canada, June 2012. To appear

References and Acknowledgments

Vaibhav Srivastava

Collaborators: R. Carli (Padova), C. Langbort (Illinois), F. Pasqualetti (UCSB), A. Surana (UTRC), Christopher Ho (UCSB), Miguel Eckstein (UCSB)

Funding: AFOSR MURI Program "Behavioral Dynamics in Mixed Human/Robotics Teams" 5/07-6/12

Attention Allocation Strategies

V. Srivastava, R. Carli, C. Langbort, and F. Bullo. Attention allocation for decision making queues. *Automatica*, February 2012. Submitted

V. Srivastava, A. Surana, and F. Bullo. Adaptive attention allocation in human-robot systems. In *American Control Conference*, pages 2767–2774, Montréal, Canada, June 2012

V. Srivastava, C. J. Ho, M. P. Eckstein, and F. Bullo. Handling operator overload: An experimental study. In preparation

Search and Surveillance Strategies

V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *International Journal of Robotics Research*, April 2012. Submitted

V. Srivastava, K. Plarre, and F. Bullo. Randomized sensor selection in sequential hypothesis testing. *IEEE Transactions on Signal Processing*, 59(5):2342–2354, 2011

General Human-Automaton Interaction Model

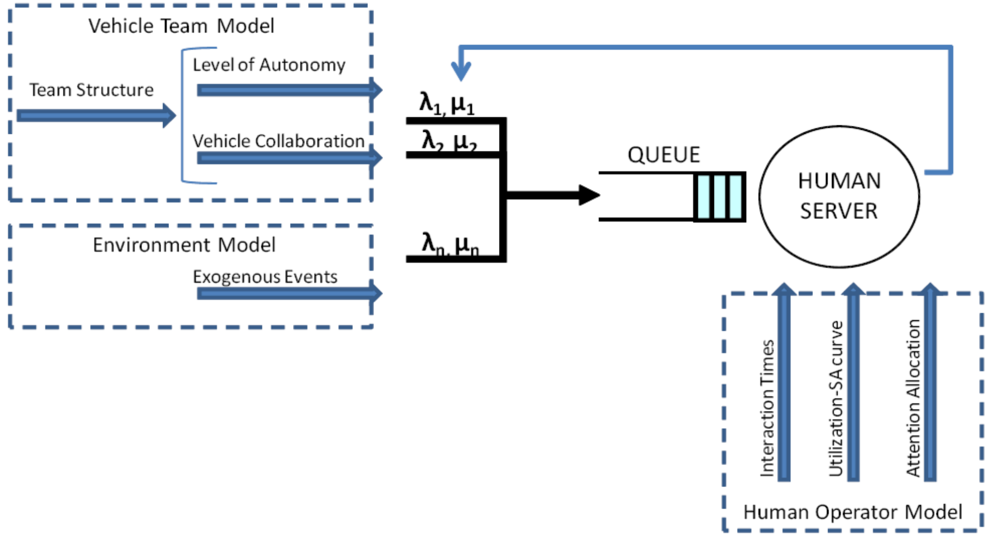

General vehicle team and human operator interaction model

C. Nehme, B. Mekdeci, J. W. Crandall, and M. L. Cummings. The impact of heterogeneity on operator performance in futuristic unmanned vehicle systems. *The International C2 Journal*, 2(2):1–30, 2008

Cognition and Autonomy Management System

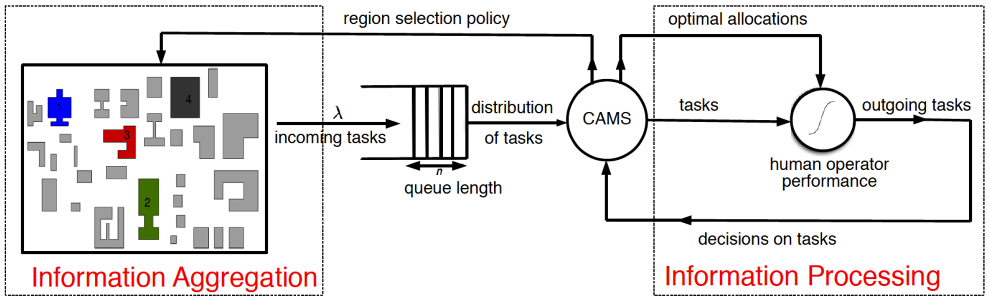

- Information collection and aggregation by robotic network
- Information processing and decision making by human operator
- Based on tasks in queue and estimated cognitive state, CAMS **specifies the time** the operator should spend on each task
- Based on the operator's decisions and world estimate, the CAMS collects information from the **most pertinent source**

Outline

Operator Cognition Models

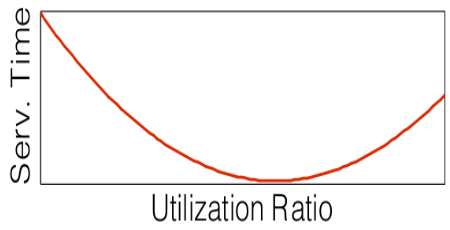

Yerkes Dodson effect

Yerkes-Dodson 1908

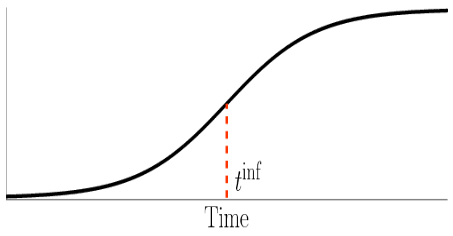

Evolution of probability of detection

Pew '68

1. operator utilization ratio = linear dynamical system
   expected (unforced) service time = convex function of utilization
   Y-D curve well-established, e.g., validated by Savla et. al. '10
2. the evidence for decision making evolves as a drift-diffusion process
   the probability of the correct decision is a sigmoid function of time

Experimental Validation of Sigmoidal Performance in Visual Perception

- task = spot the differences
- expected # detected differences
  is linear function of time (DDM)
- probability to detect more than 60% diffs
  is sigmoid (threshold-based decision making)

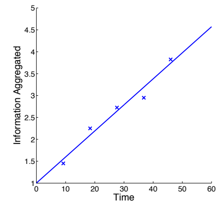

Information aggregation satisfy DDM

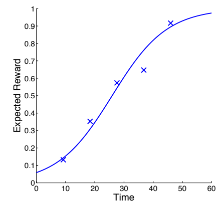

Probability of correct decision is sigmoid

Implications of Sigmoid Performance

Sigmoid function and linear penalty

$$\underset{t \geq 0}{\text{maximize}} \quad f(t) - \psi t$$

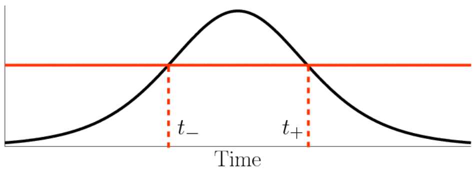

Derivative of a sigmoid function

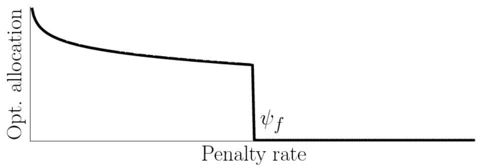

Optimal allocation v/s penalty rate

- **Optimal allocation jumps down to zero at critical penalty rate**
- **Jump creates combinatorial effects**

Dynamic Queue with Penalty and Situational Awareness I

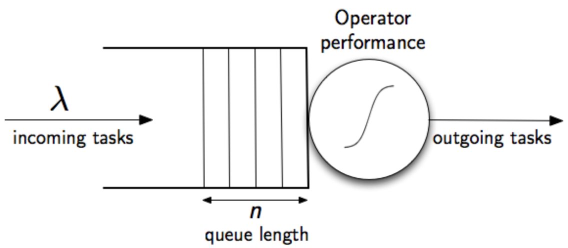

- Tasks arrive as a Poisson process with rate λ
- Task γ sampled from a distribution
  reward w_γ, sigmoid params (inflection, slope), penalty rate c_γ
- **State variables:** queue length n_ℓ and utilization ratio x_ℓ at stage ℓ
- Unforced service time = Y-D law $S_\gamma(x)$
- **Decision variables:** duration allocation t_ℓ, rest time r_ℓ, binary z_ℓ

Dynamic Queue with Penalty and Situational Awareness II

Average Reward

$$\max_{z_\ell t_\ell \geq z_\ell S(x_{\ell-1})} \lim_{L\to\infty} \frac{1}{L}\sum_{\ell=1}^{L} z_\ell\Big(\mathbb{E}[w_{\gamma_\ell} f_{\gamma_\ell}(t_\ell)] - \bar{c}\,\mathbb{E}[n_\ell](t_\ell+r_\ell) - \frac{\bar{c}\lambda(t_\ell+r_\ell)^2}{2}\Big)$$

System Dynamics

Queue length: $\mathbb{E}[n_{\ell+1}] = \mathbb{E}[\max\{1, n_\ell - 1 + \text{Poisson}(\lambda z_\ell(t_\ell + r_\ell))\}]$,

Utilization: $x_{\ell+1} = (1 - e^{-\frac{t_\ell z_\ell}{\tau}} + x(\ell)e^{-\frac{t_\ell z_\ell}{\tau}})e^{-\frac{r_\ell z_\ell}{\tau}}, \; x_\ell \in [x_{\min}, x_{\max}]$

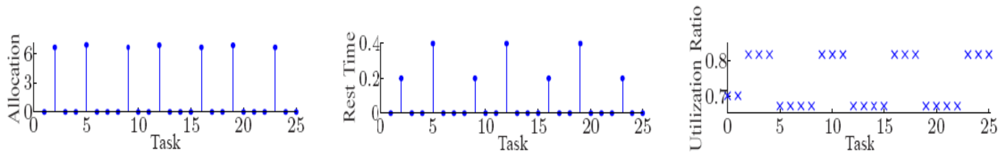

Certainty Equivalent Solution

Illustrative Example I

Optimal Allocations and Rest Time

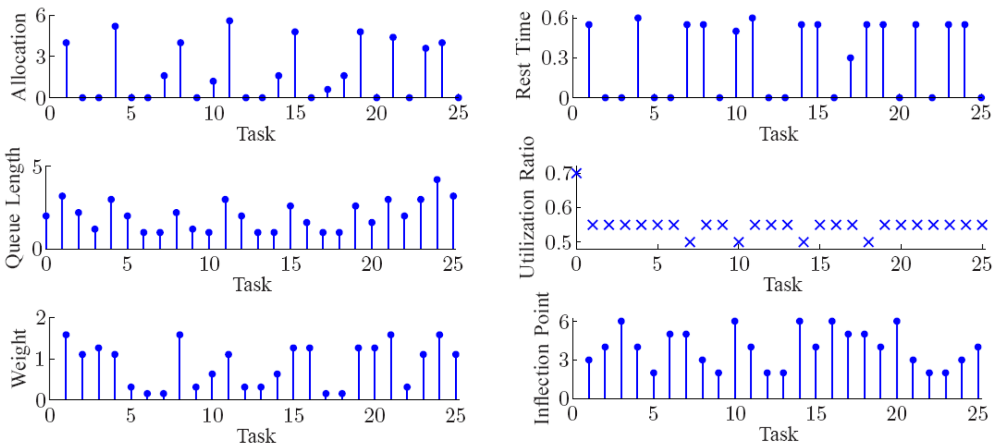

Receding Horizon Policy

Illustrative Example II

Reward versus Arrival Rate

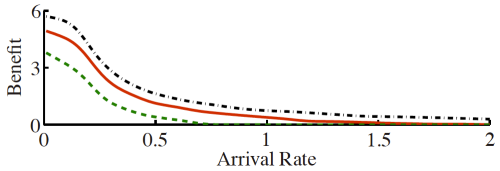

Benefit per unit task

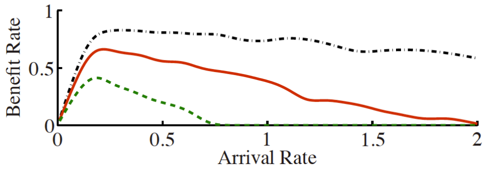

Benefit rate

Optimal arrival rate

- Switching occurs when operator is expected to be always non-idle
- Designer may pick desired accuracy on each task to design arrival rate

Outline

1. Introduction
2. Topic 1: Information Processing
3. Topic 2: Information Aggregation
4. Topic 3: Combined Information Aggregation and Processing
5. Conclusions

Cognition and Autonomy Management System

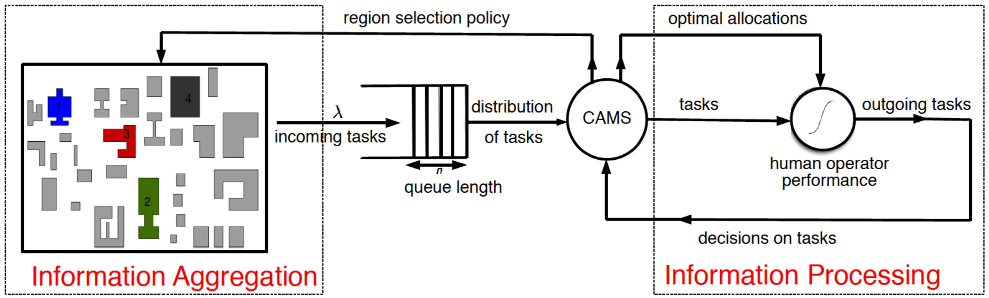

- Information collection and aggregation by robotic network
- Information processing and decision making by human operator

- Based on tasks in queue and estimated cognitive state, CAMS **specifies the time** the operator should spend on each task
- Based on the operator's decisions and world estimate, the CAMS collects information from the **most pertinent source**

Spatial Quickest Detection

Dynamic Vehicle Routing for Distributed Surveillance

- N regions, arbitrary # anomalies
- an ensemble of CUSUM algorithms
- T_k = collection + transmission + processing time at region k
- d_{ij} = distance between region i and j

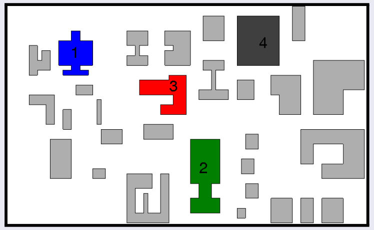

UCSB Campus Map

Spatial Quickest Detection

1. at iteration τ, pick a region k from stationary distribution $\mathbf{q}$
2. go to region k and collect evidence y_τ
3. update CUSUM statistic for region k

$$\Lambda_k = (\Lambda_{k-1} + \log(f_k^1(y_\tau)/f_k^0(y_\tau)))^+$$

4. declare an anomaly at region k if $\Lambda_k > \eta$

Spatial Quickest Detection: Detection Delay

Expected detection delay at region k

$$\mathbb{E}[\text{delay}_k(\mathbf{q})] = \frac{e^{-\eta} + \eta - 1}{q_k \mathcal{D}(f_k^1, f_k^0)}(\mathbf{q} \cdot \mathbf{T} + \mathbf{q} \cdot D\mathbf{q})$$

Two stage quickest detection strategy

1. pick optimal $\mathbf{q}^* = \text{argmin} \sum_{k=1}^{N} \pi_k^1 \, \mathbb{E}[\text{delay}_k(\mathbf{q})]$
2. adapt $\mathbf{q}^*$ with the evidence collected at each stage

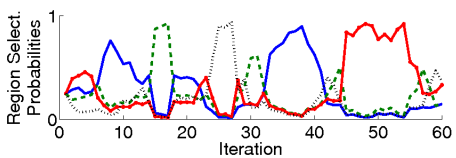

Region Selection Probability

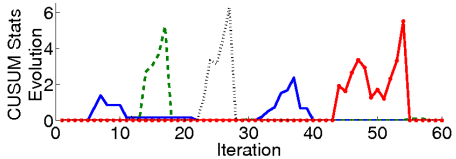

Likelihood of Anomaly

Outline

Cognition and Autonomy Management System

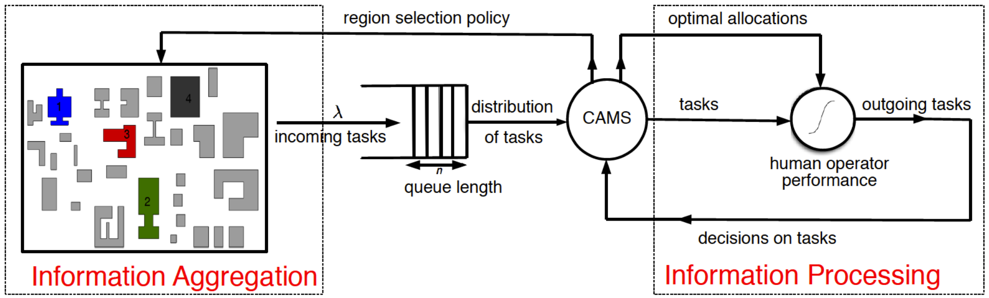

- Information collection and aggregation by robotic network
- Information processing and decision making by human operator

- Based on tasks in queue and estimated cognitive state, CAMS **specifies the time** the operator should spend on each task
- Based on the operator's decisions and world estimate, the CAMS collects information from the **most pertinent source**

Spatial Quickest Detection with Human Input

- human operator allocates time t to a task and decides on presence/absence of anomaly
- decision in a Bernoulli random variable with

$$\mathbb{P}(\text{success}|t) = \begin{cases} f_k^1(t), & \text{if an anomaly is present,} \\ f_k^0(t), & \text{if no anomaly is present.} \end{cases}$$

Spatial Quickest Detection

1. at stage ℓ, pick a region k from stationary distribution $\mathbf{q}$
2. go to region k and collect evidence y_ℓ and decision $\text{dec}_\ell \in \{0,1\}$
3. update CUSUM statistic for region k

$$\Lambda_k = (\Lambda_{k-1} + \log(\mathbb{P}(\text{dec}_\ell | t_\ell, \text{anomaly}) / \mathbb{P}(\text{dec}_\ell | t_\ell, \text{no anomaly}))^+$$

4. declare an anomaly at region k if $\Lambda_k > \eta$

Simultaneous Information Aggregation and Processing I

Critical Issue:
human decisions are not i.i.d.
$\implies$ **no closed form delay expression**

Upper Bound

$$\mathbb{E}[\text{delay}_k] \leq \frac{e^{-\eta}+\eta-1}{q_k \mathcal{D}_{\min}(k)}(\mathbf{q}\cdot\mathbf{T}+\mathbf{q}\cdot D\mathbf{q})$$

Adaptive Policy with Human Feedback

1. determine $\mathbf{q}^*$ and sample regions
2. set operator performance at region k

$$f_k(t) = \pi_k f_k^1(t) + (1-\pi_k) f_k^0(t)$$

3. determine optimal allocation and rest time
4. update CUSUM statistic using operator's decision
5. go to step 1.

Simultaneous Information Aggregation and Processing II

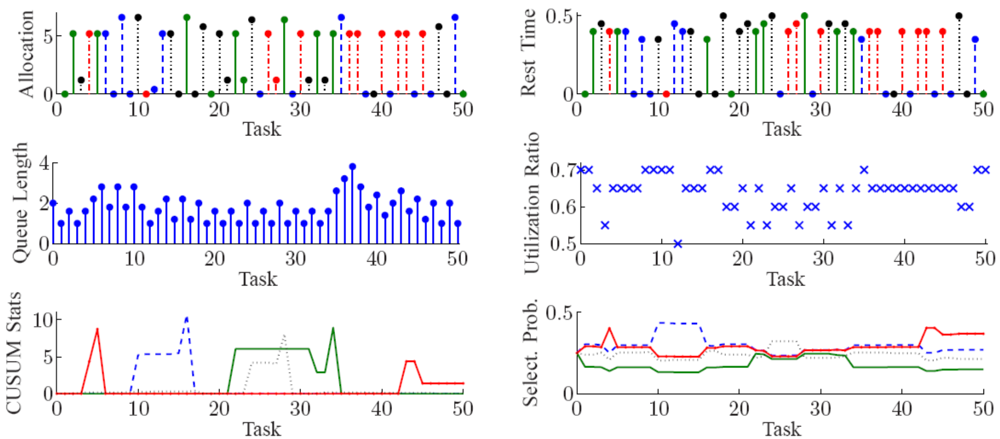

Optimal Policies

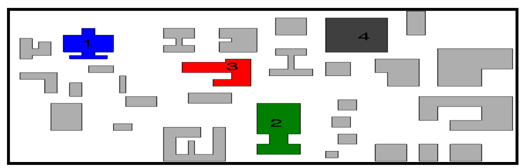

Outline

Conclusions & Future Directions

Conclusions

- disciplines: human cognitive performance models, dynamics vehicle routing, decision making, dynamic optimization
- *simultaneous information aggregation and processing* architecture
- incorporation of cognitive / situational awareness / autonomy
- an adaptive strategy that collects evidence from regions with high likelihood of anomalies and optimally processes it

Future Directions

- experimental validation of models [ongoing] and of architecture
- incorporation of fatigue, learning and other cognitive models
- re-queuing of tasks, preemptive queues and more general scenarios
- dynamic anomalies and more complex detection tasks
- multi-vehicle, multi-operator, single-operator multitasking, heterogeneous scenarios

3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems

NecSys'12, September 14-15, 2012, Fess Parker's Doubletree Resort, Santa Barbara, California

Relevant Dates and Proceedings

- Submissions to NecSys 12 are open as of March 25. Please, read the Information for Authors.
- Extended Papers submission deadline: April 30, 2012
- Notice of acceptance: June 14, 2012
- Final version due: July 15, 2012
- Early registration deadline: July 15, 2012
- Hotel registration deadline: August 13, 2012