Multi-robot Exploration with Limited Communication in the RoboCup Rescue

Arnoud Visser

Lisbon Workshop on New Challenges for Cooperative Robotics,
Lisbon, Portugal, October 24, 2008

Universiteit van Amsterdam
Informatica Instituut
intelligent autonomous systems

Virtual Rescue League

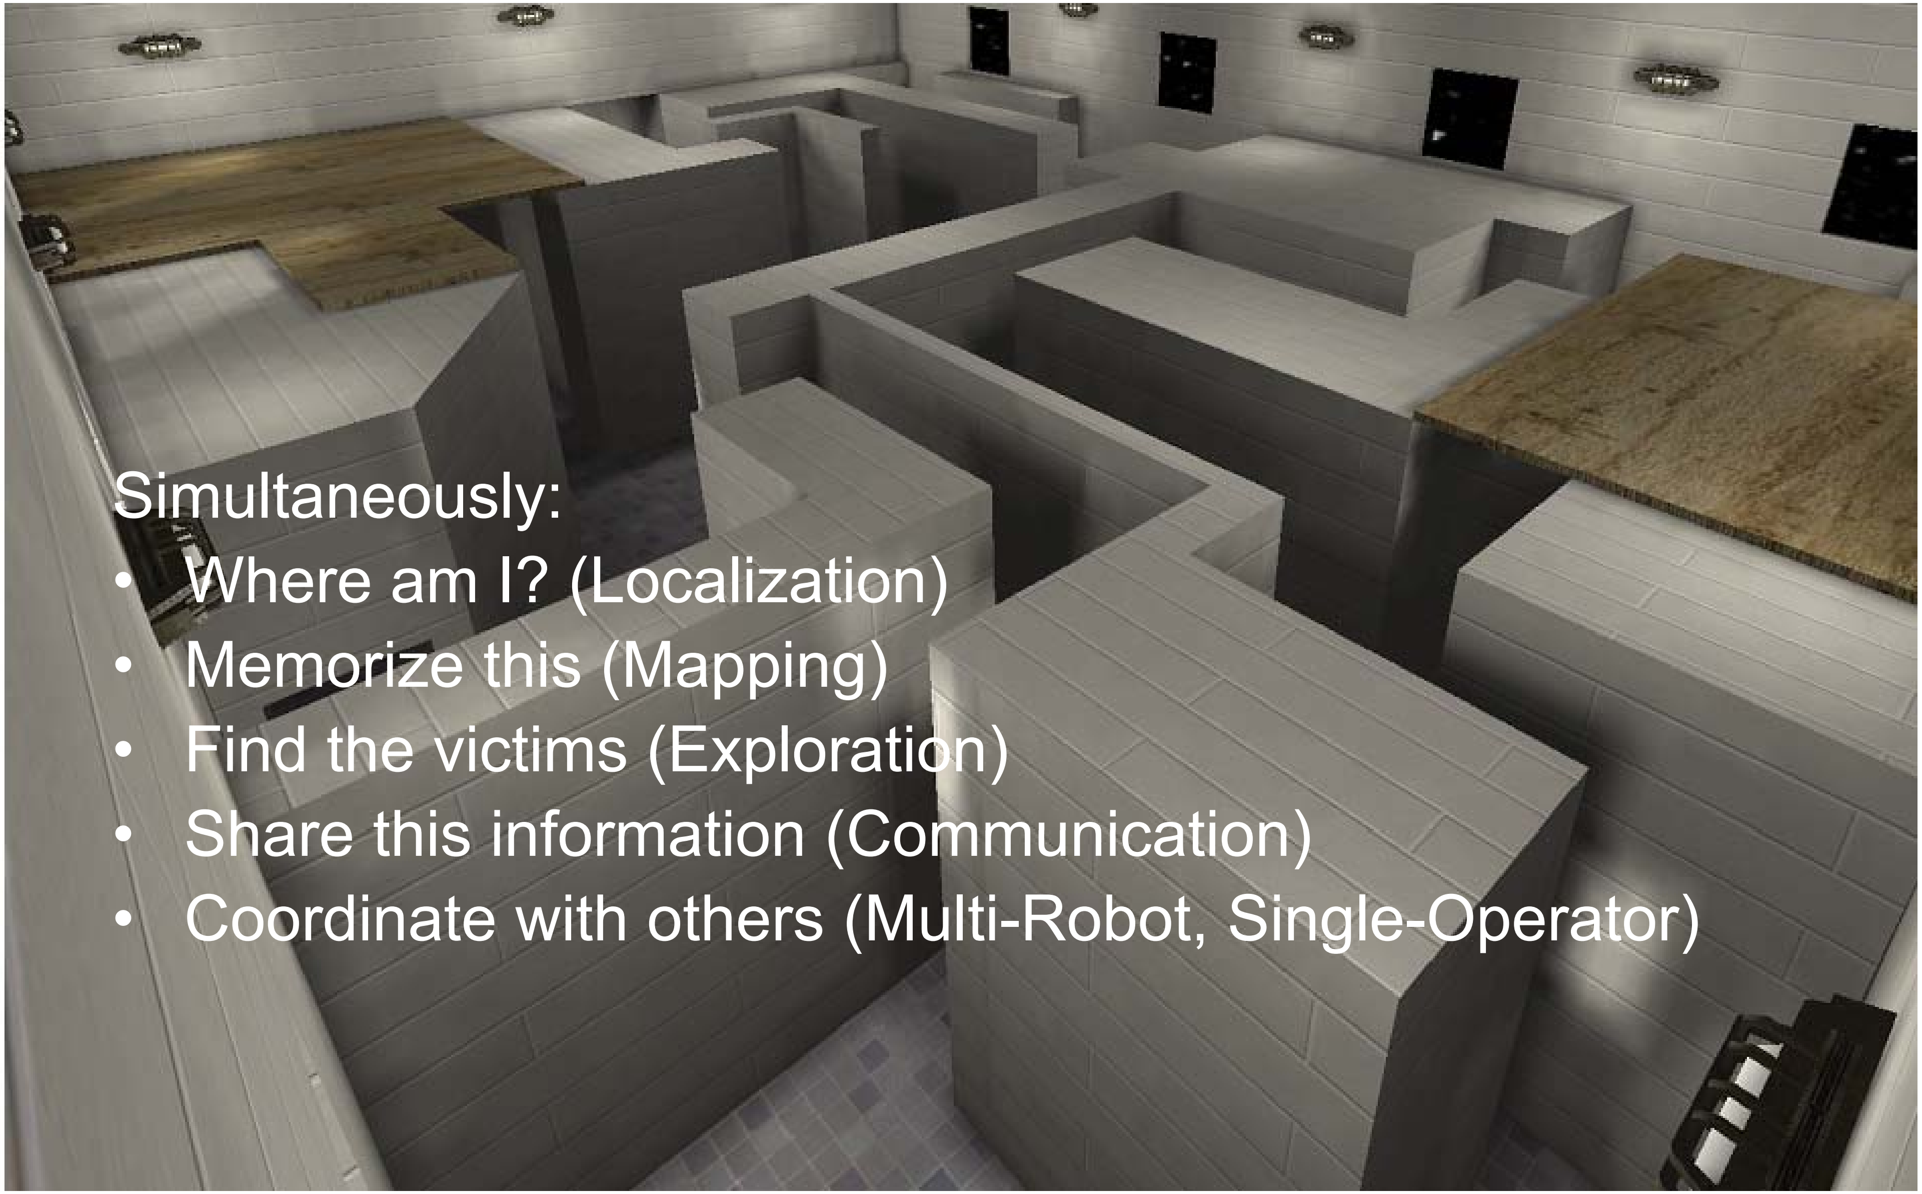

A wide variety of simulated worlds

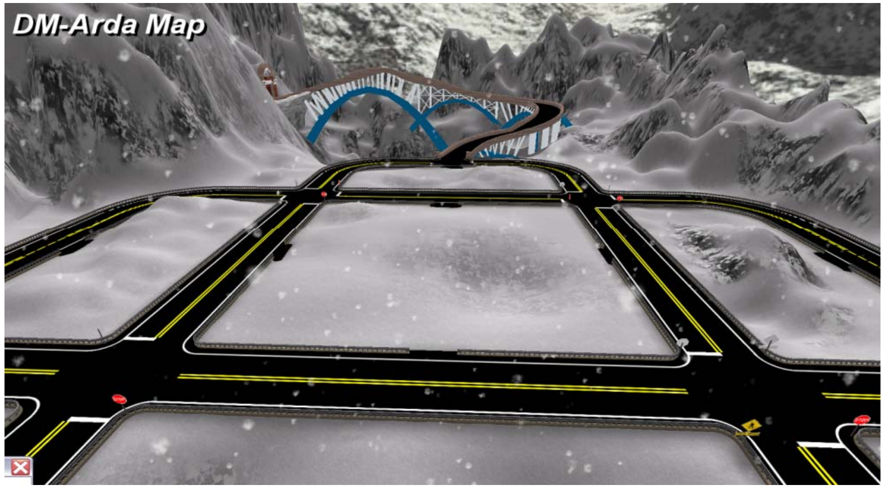

A wide variety of Robotic platforms

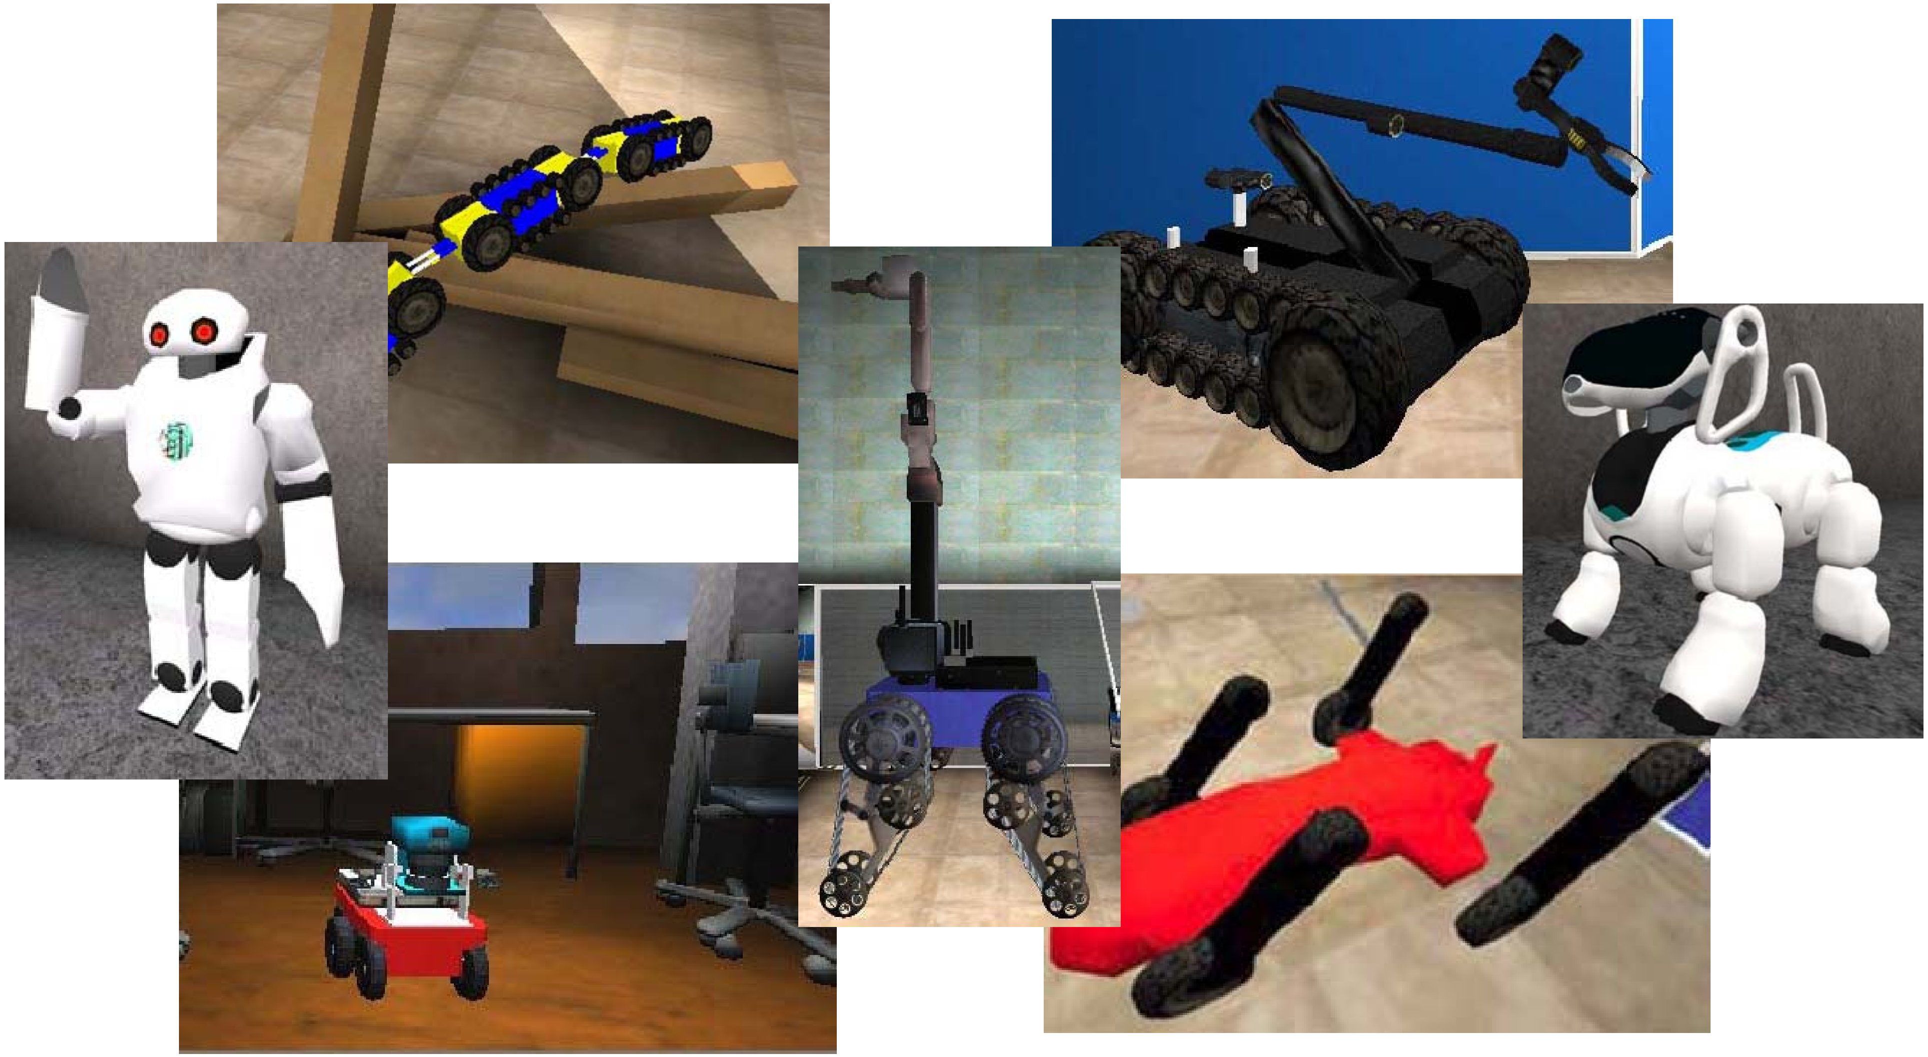

Networked robot team

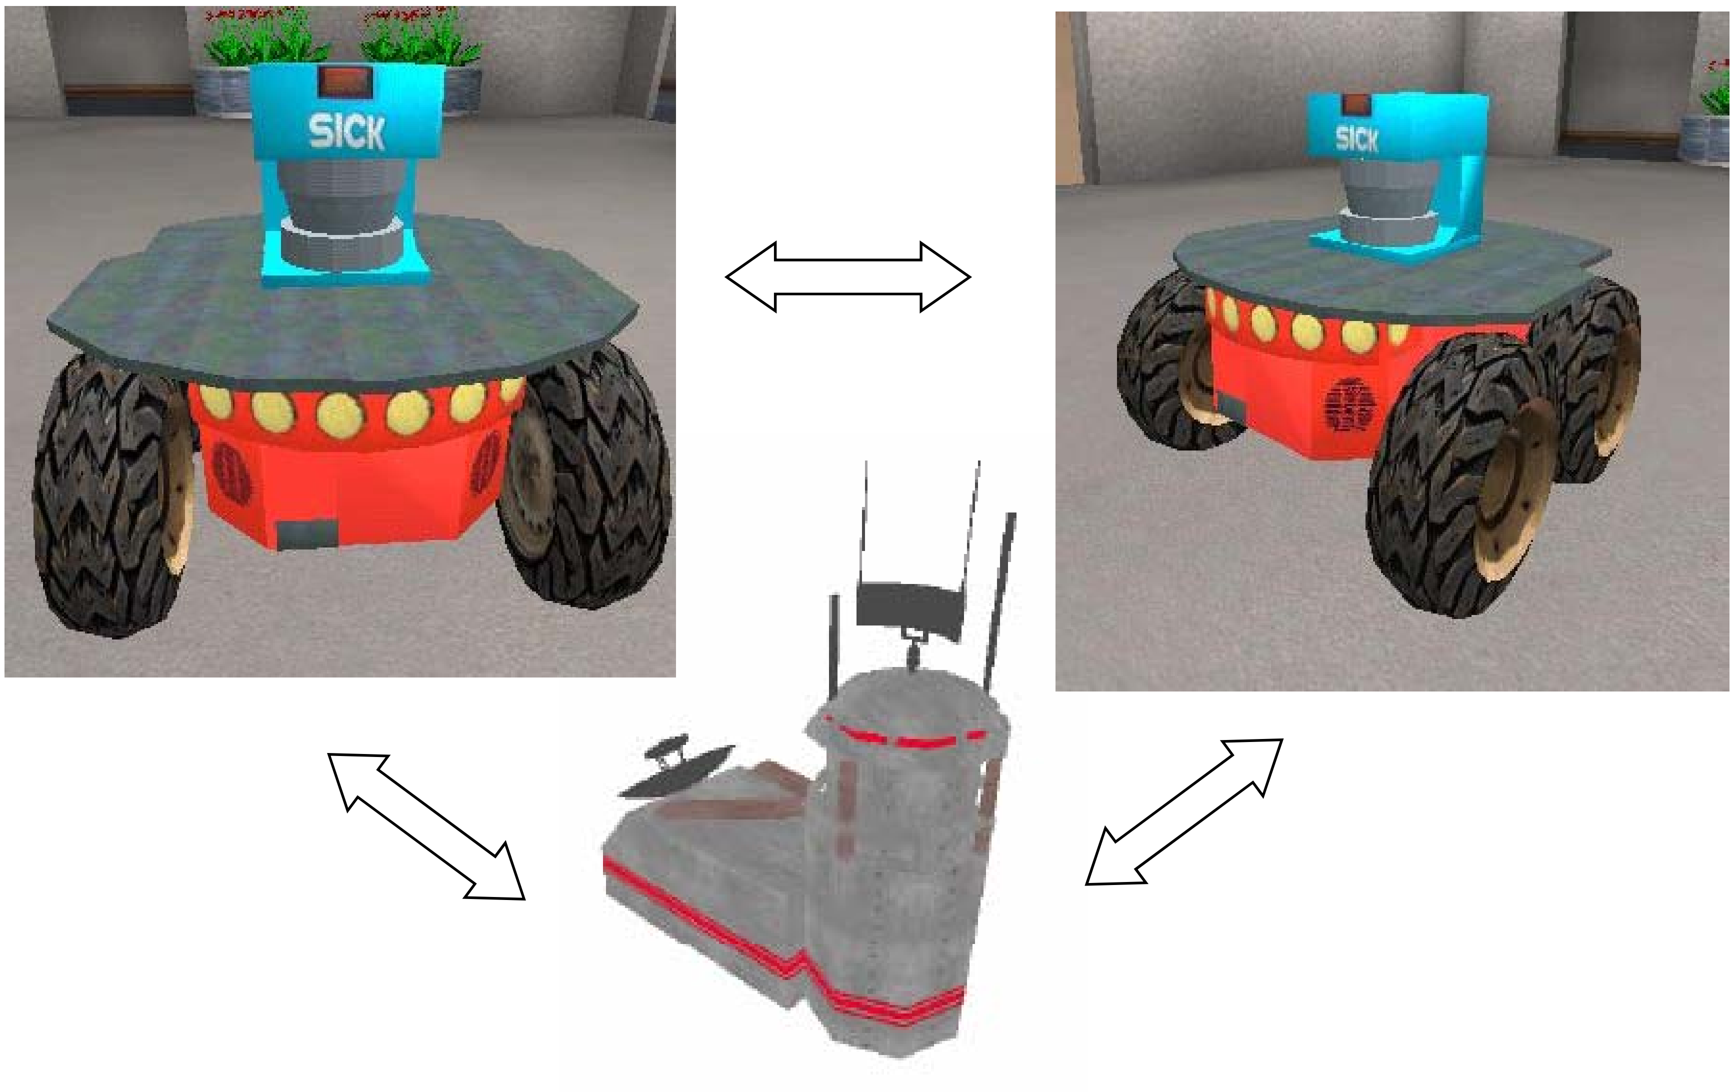

UsarCommander

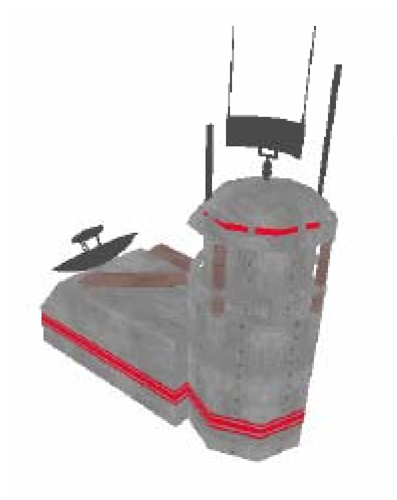

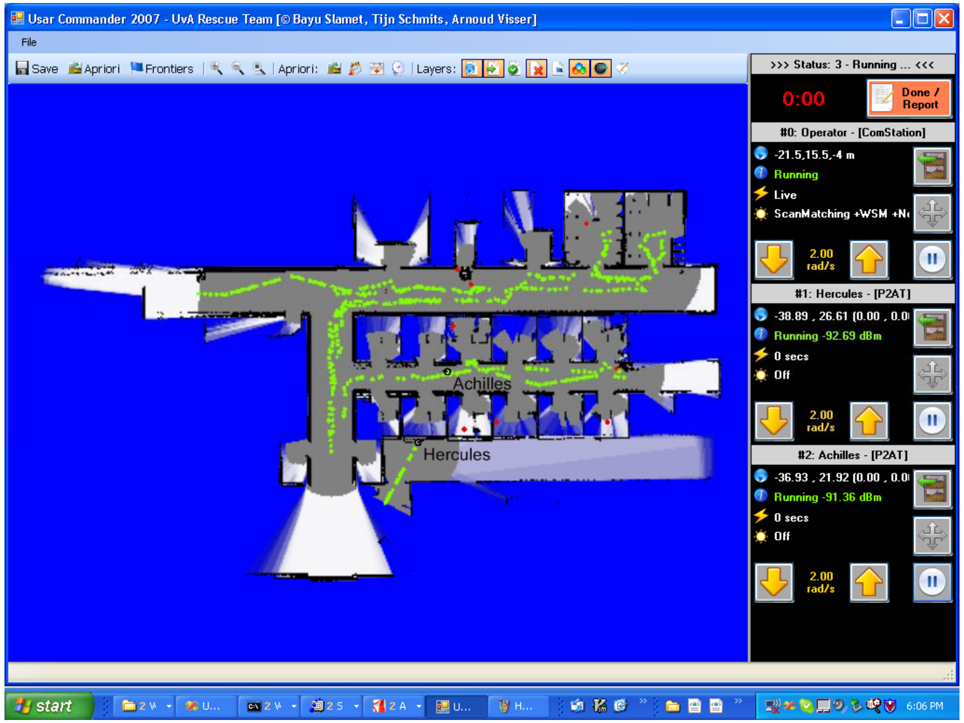

- User Interface can be used to teleoperate the robots and monitor the shared map.
- Here, the robots are autonomous exploring on a distributed map (each with a partial view).

Localization and Mapping challenge:

Combination of Manifold SLAM with Weighted Scan Matching algorithm

Bayu Slamet and Max Pfingsthorn

intelligent autonomous systems

Universiteit van Amsterdam
Informatica Instituut

Pose estimation

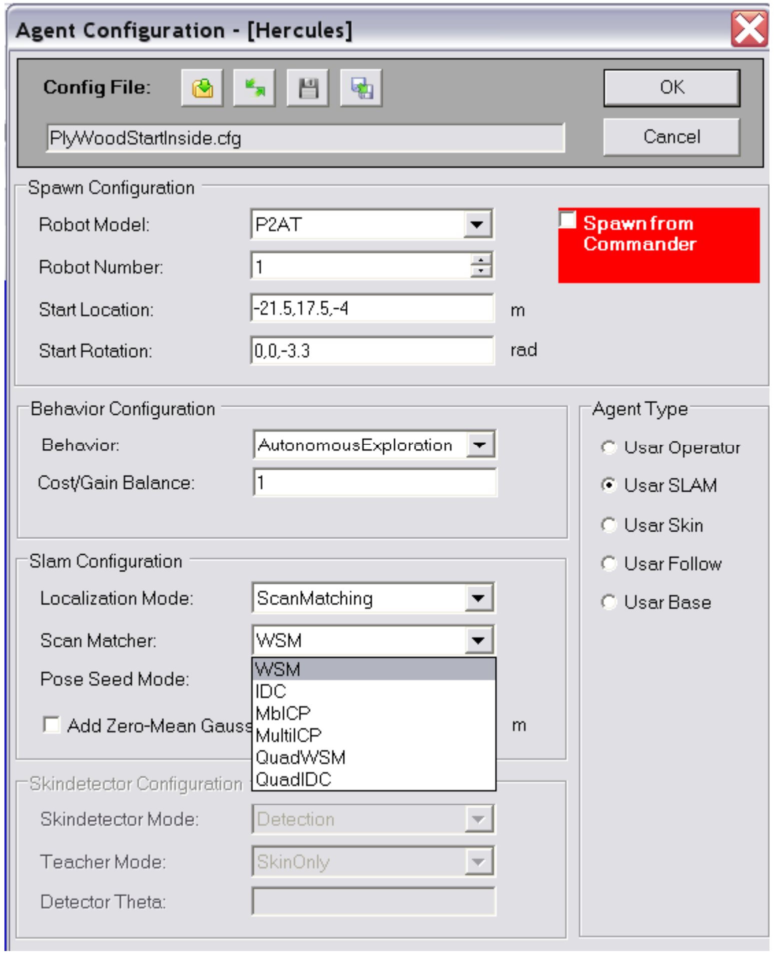

Pose & confidence

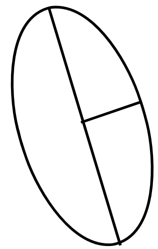

- Several Scan Matching algorithms

Localization vs Mapping

- When the confidence is good, the current position fits to one of the positions of the local submap
  (localization, no extension of the map)

- When confidence drops, a new *patch* is added to the chain
  (*graph based mapping)*

Chain of patches

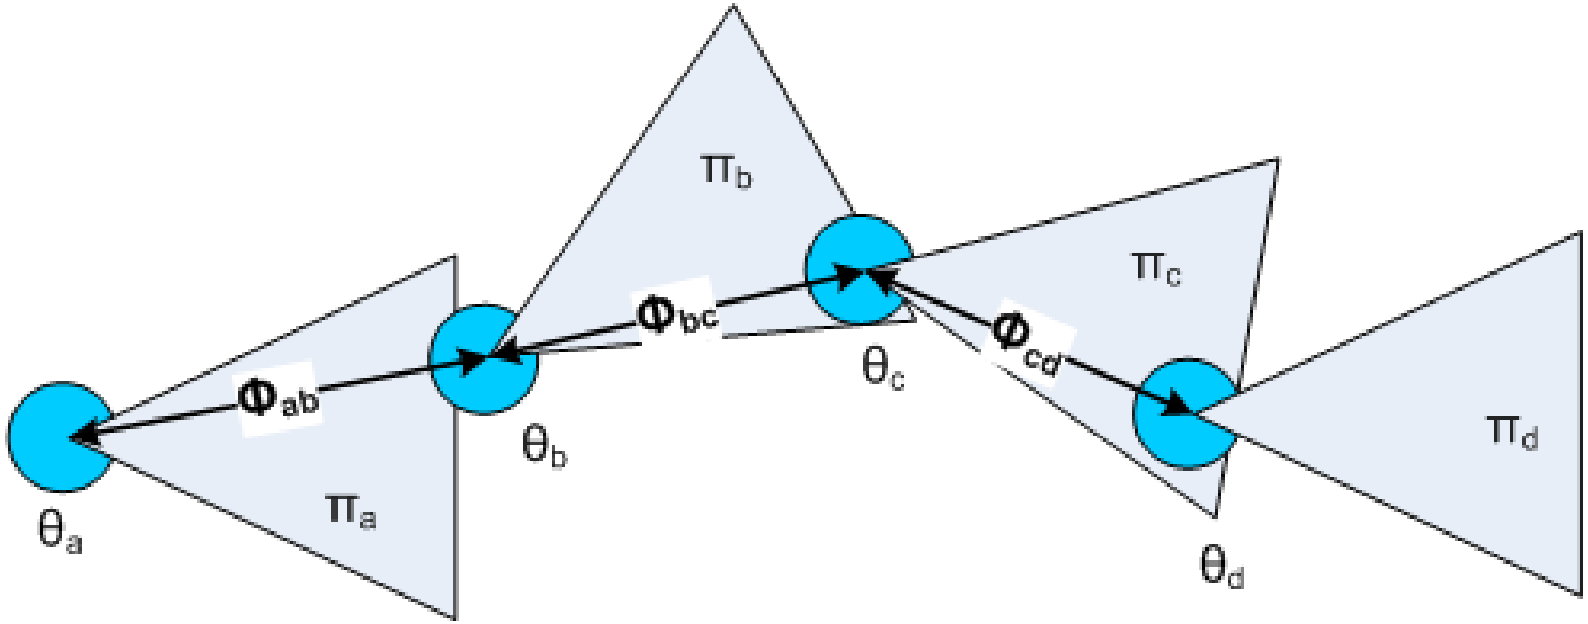

π = laser scan

θ = absolute location (Euclidean)

Φ= relation

= relative location (polar) $\Delta\theta$ + covariance matrix Σ

Island merging

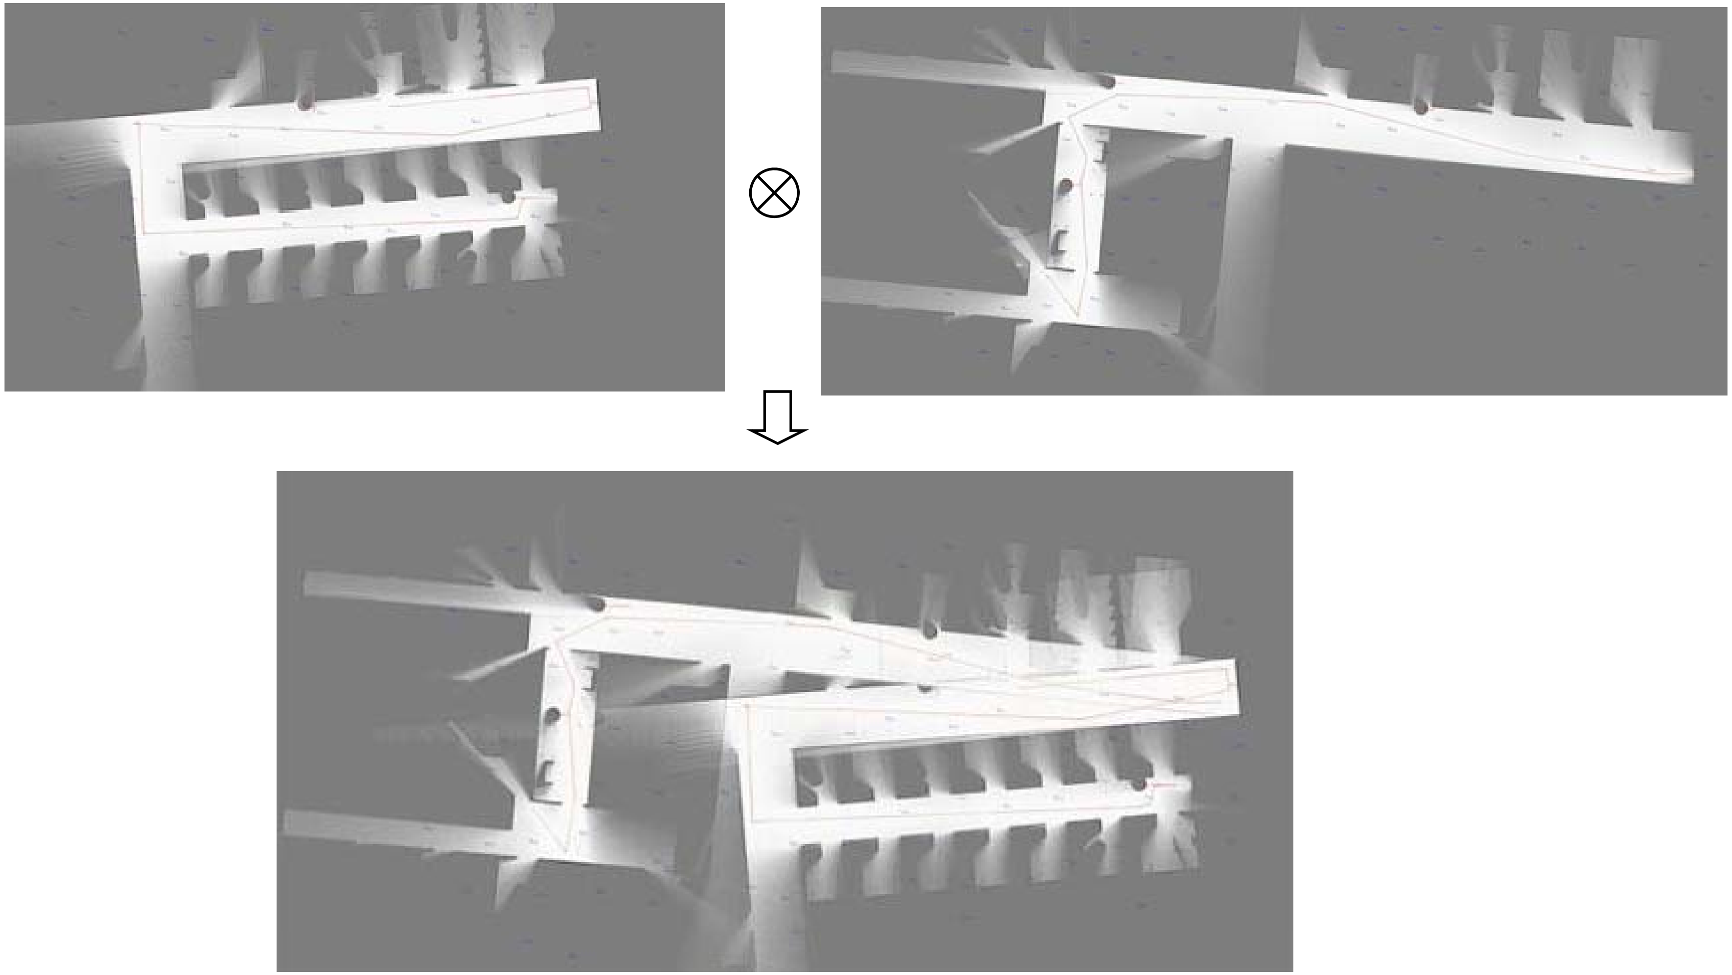

Rigid body transformation

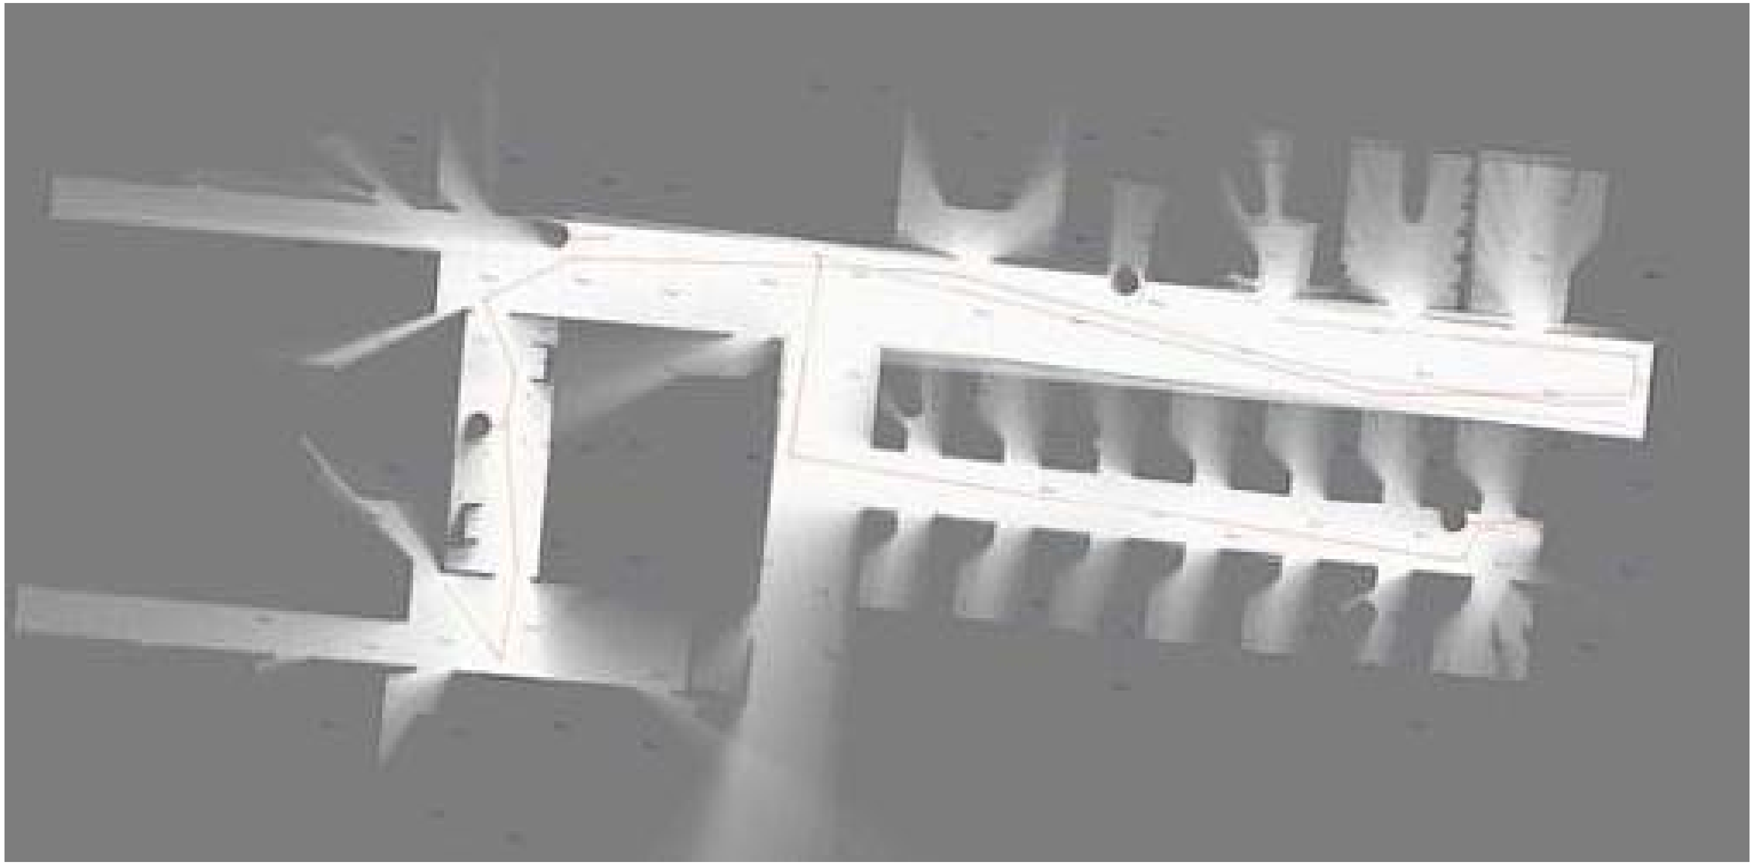

Refitting graph-relations

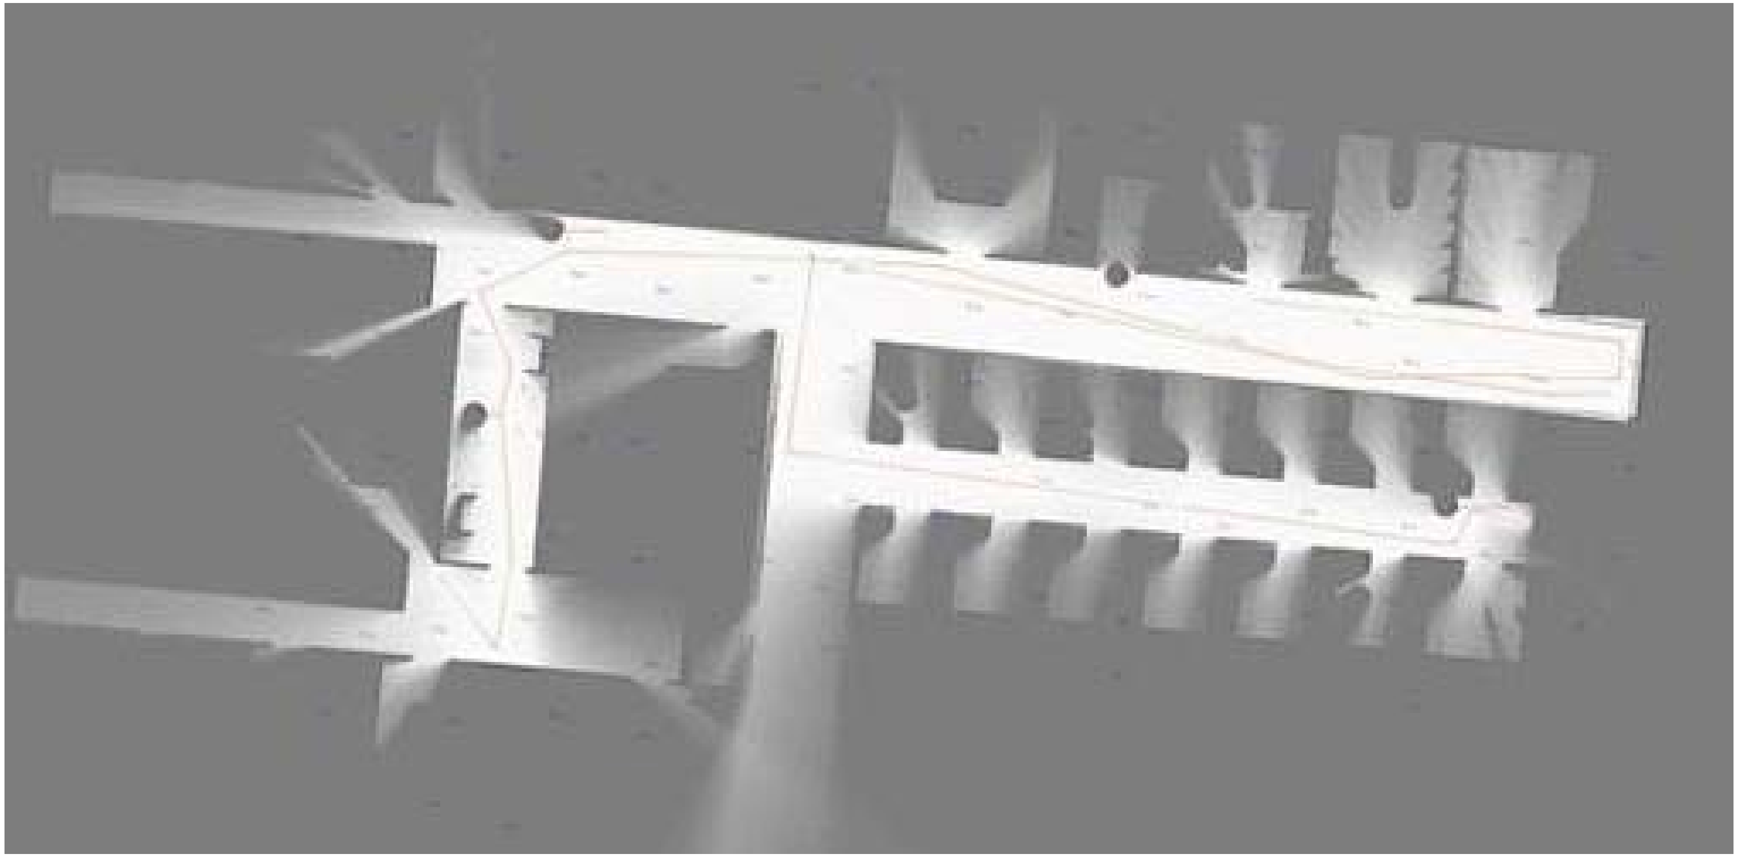

Corridor

Localization and Mapping challenge:
Result

- Distributed approach scales well for multiple robots
- High accuracy maps can be achieved
- Additional data can be added asynchronously
- Results have been validated

See Max Pfingsthorn, Bayu Slamet and Arnoud Visser, "A Scalable Hybrid Multi-Robot SLAM Method for Highly Detailed Maps", Lecture Notes on Artificial Intelligence series.

Validation on Radish datasets

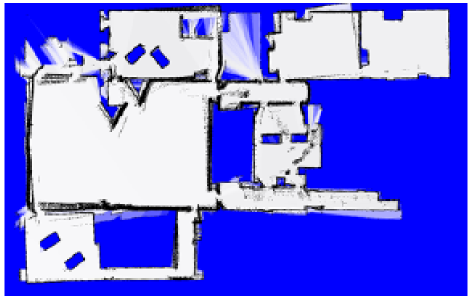

(a) AP Hill

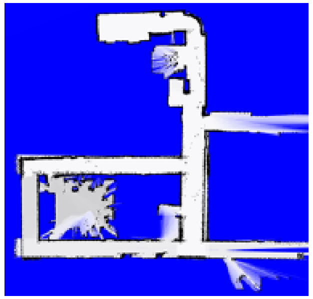

(b) CMU N. Simon Hall

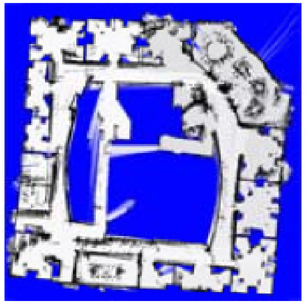

(e) Intel Lab, Seattle

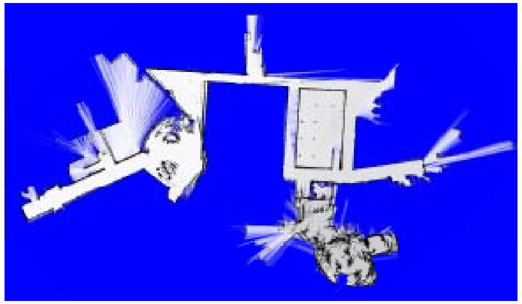

(f) 3rd floor of MIT CSAIL

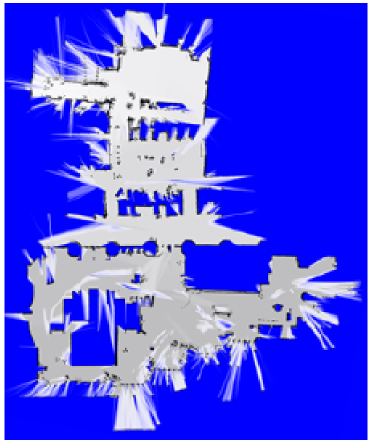

(c) Edmonton Convention

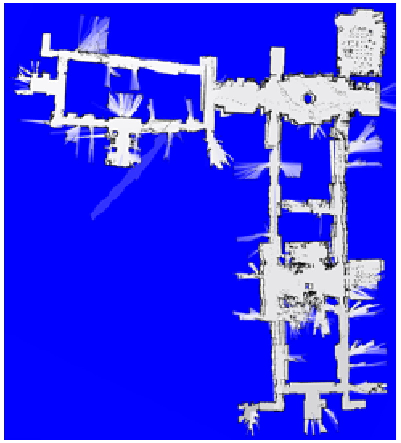

(d) Stanford Gates

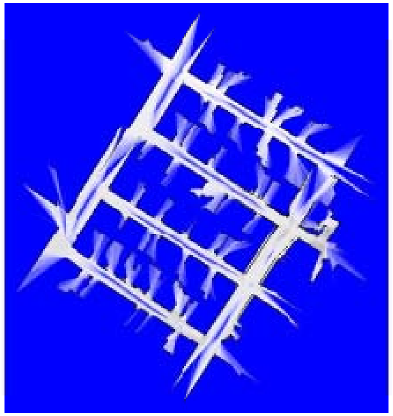

(g) Intel Campus, Oregon

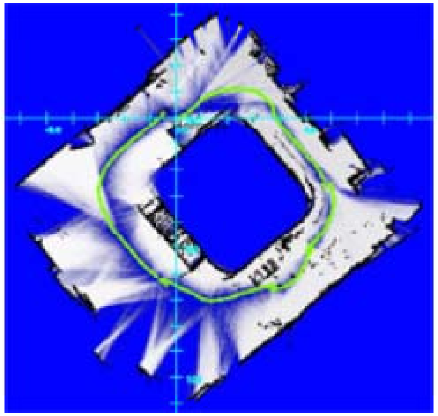

(h) NIKHEF Amsterdam

Navigation Challenge:
Beyond frontier exploration

Merlijn van Ittersum, Xingrui-Ji,
Luis Gonzalez and Laurentiu Stancu

intelligent autonomous systems

Universiteit van Amsterdam
Informatica Instituut

Interpret the map

- Make the information on the map available in a value function to optimize the selection of the next action.

- The value function has to balance the information gain against the cost of the action.

Information gain

- Active localization can be seen as minimizing the entropy *H(m)* of map *m*

$$H(m) = -\sum_{x \in m} p(x) log(p(x))$$

- Important is the difference in entropy *H* before and after an exploration action *a*

$$\Delta I(a) = H(m|a) - H(m)$$

Reduction of action space

- At a large map many exploration actions are possible. A finite number of observation points have to be chosen.
- Typically, those observation points are chosen on the boundaries of explored and unexplored areas: frontier based exploration

See Yamauchi, B.: A frontier based approach for autonomous exploration. In: Proceedings of IEEE International Symposium on Computational Intelligence in Robotics and Automation, Monterey, CA, July 10-11, 1997. (1997)

Candidate Observation points

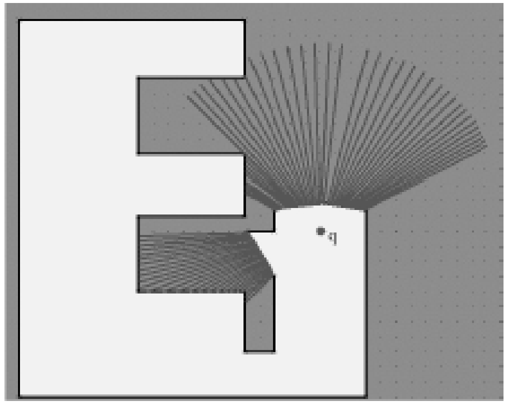

Courtesy of H.H. Gonzalez-Banos et al.

- The potential information gain *ΔI(a)* of a candidate observation point *q* is the area *A(q)* that may be visible through the two free edges; this area is estimated by casting rays from *q*.

Probabilistic occupancy map

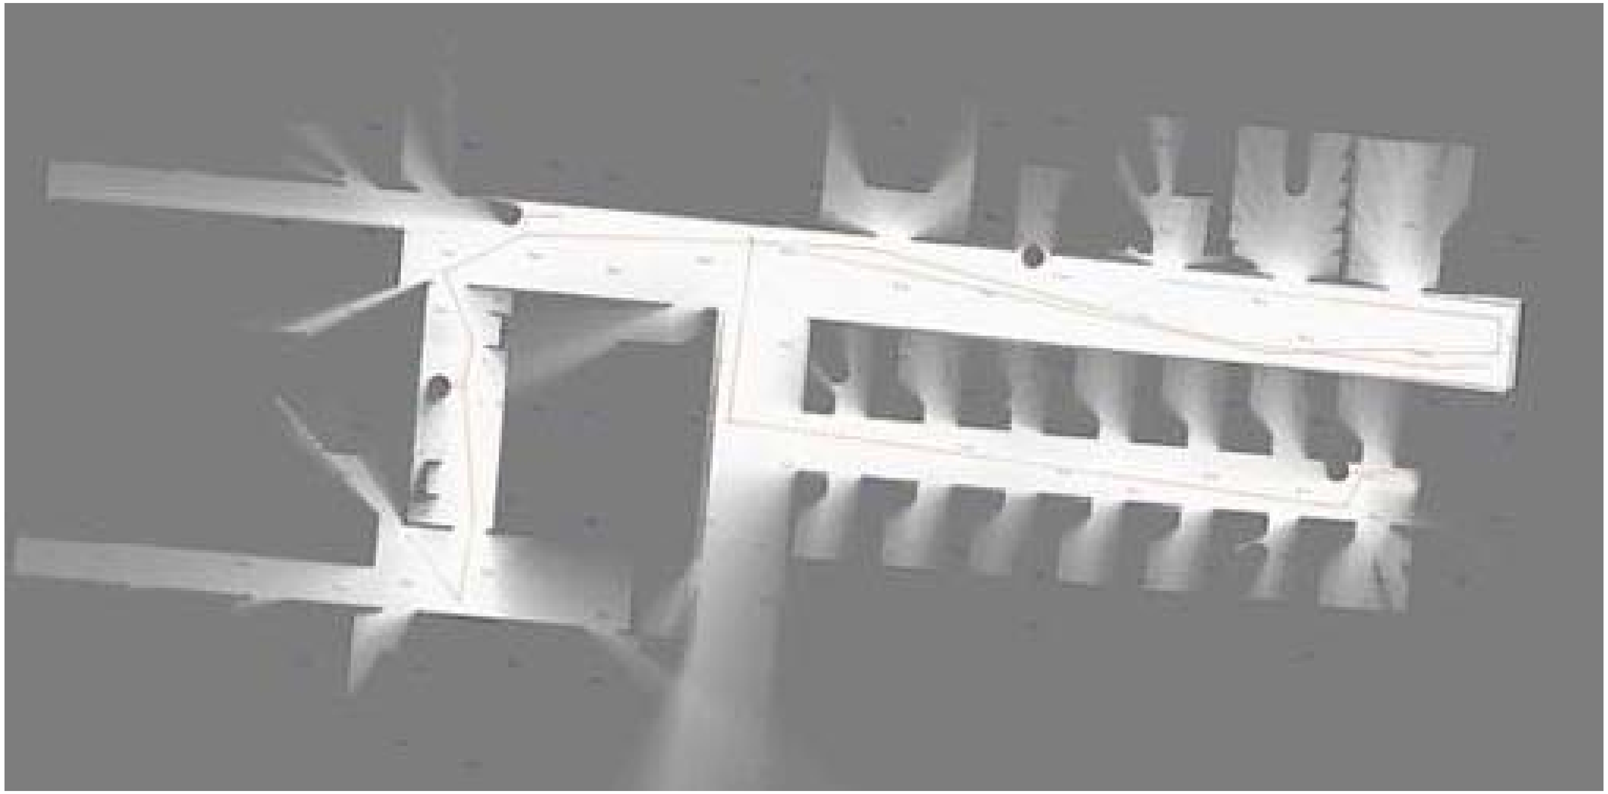

Ray-tracing with a long and short maximum range

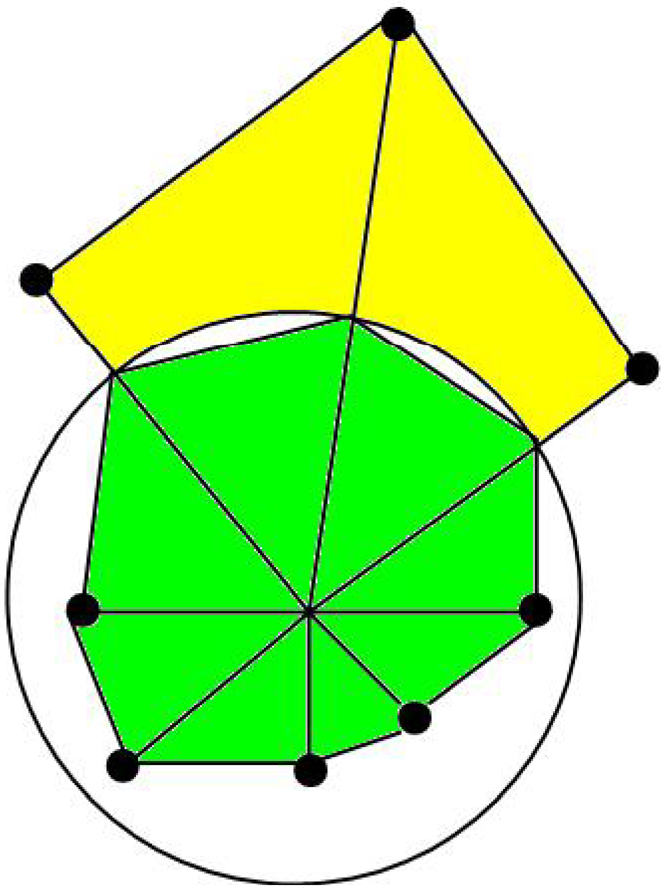

How to interpret the map

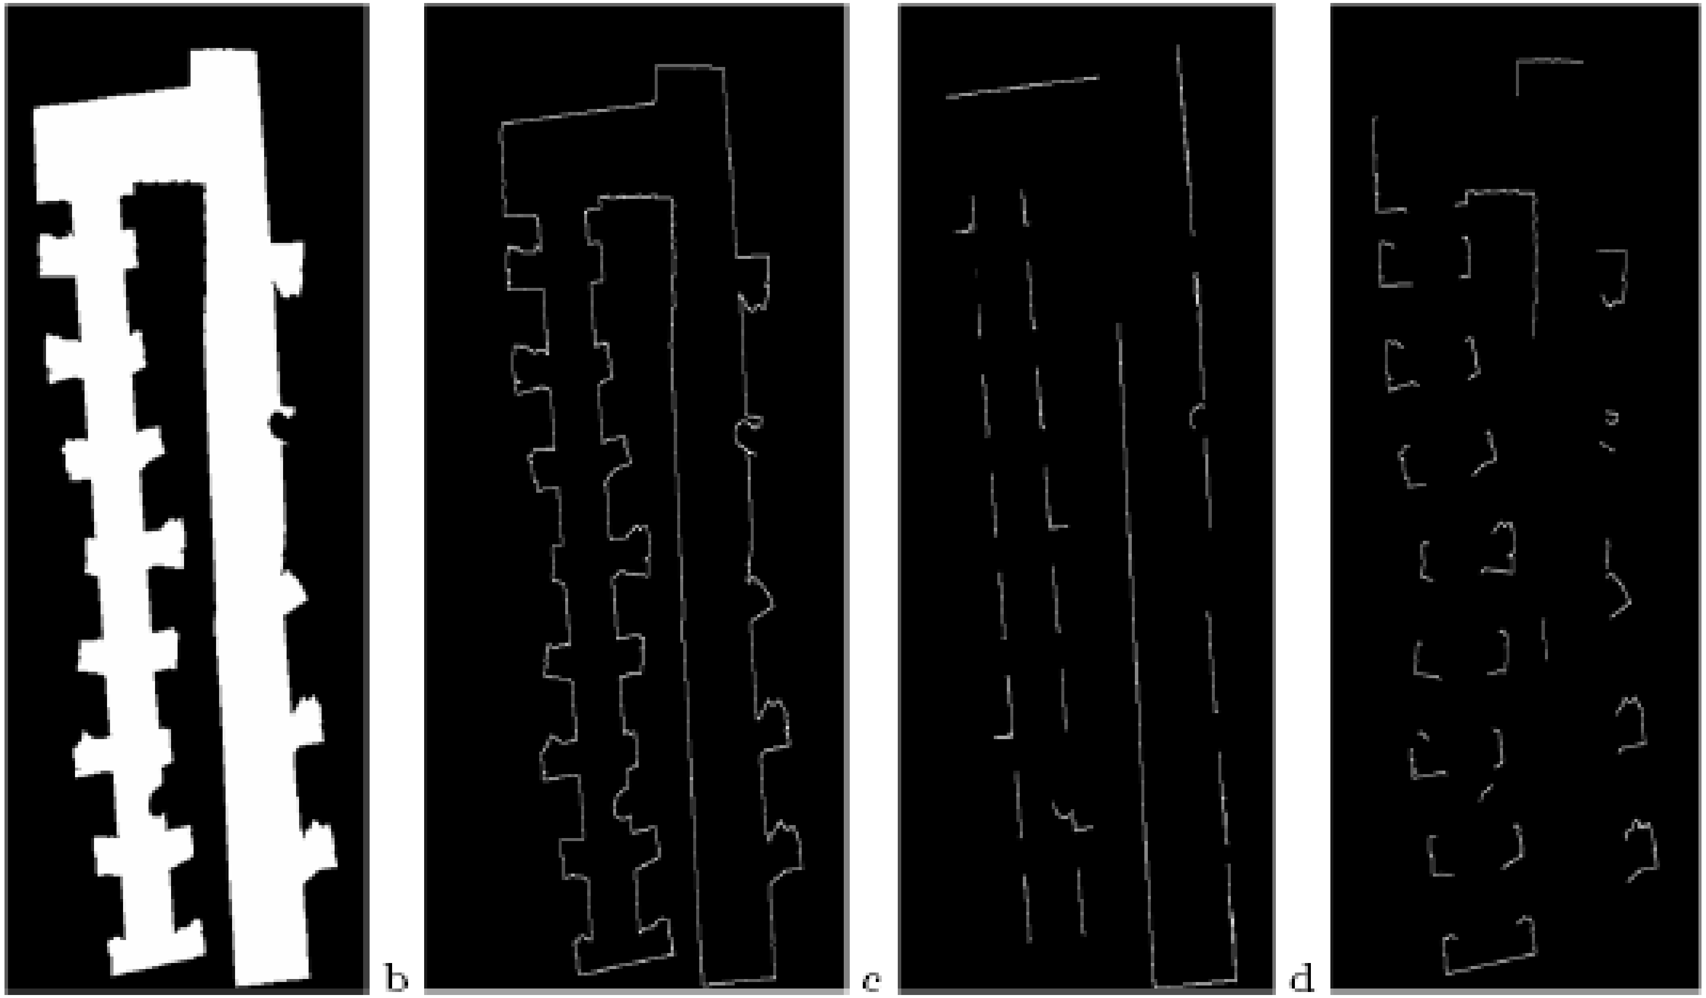

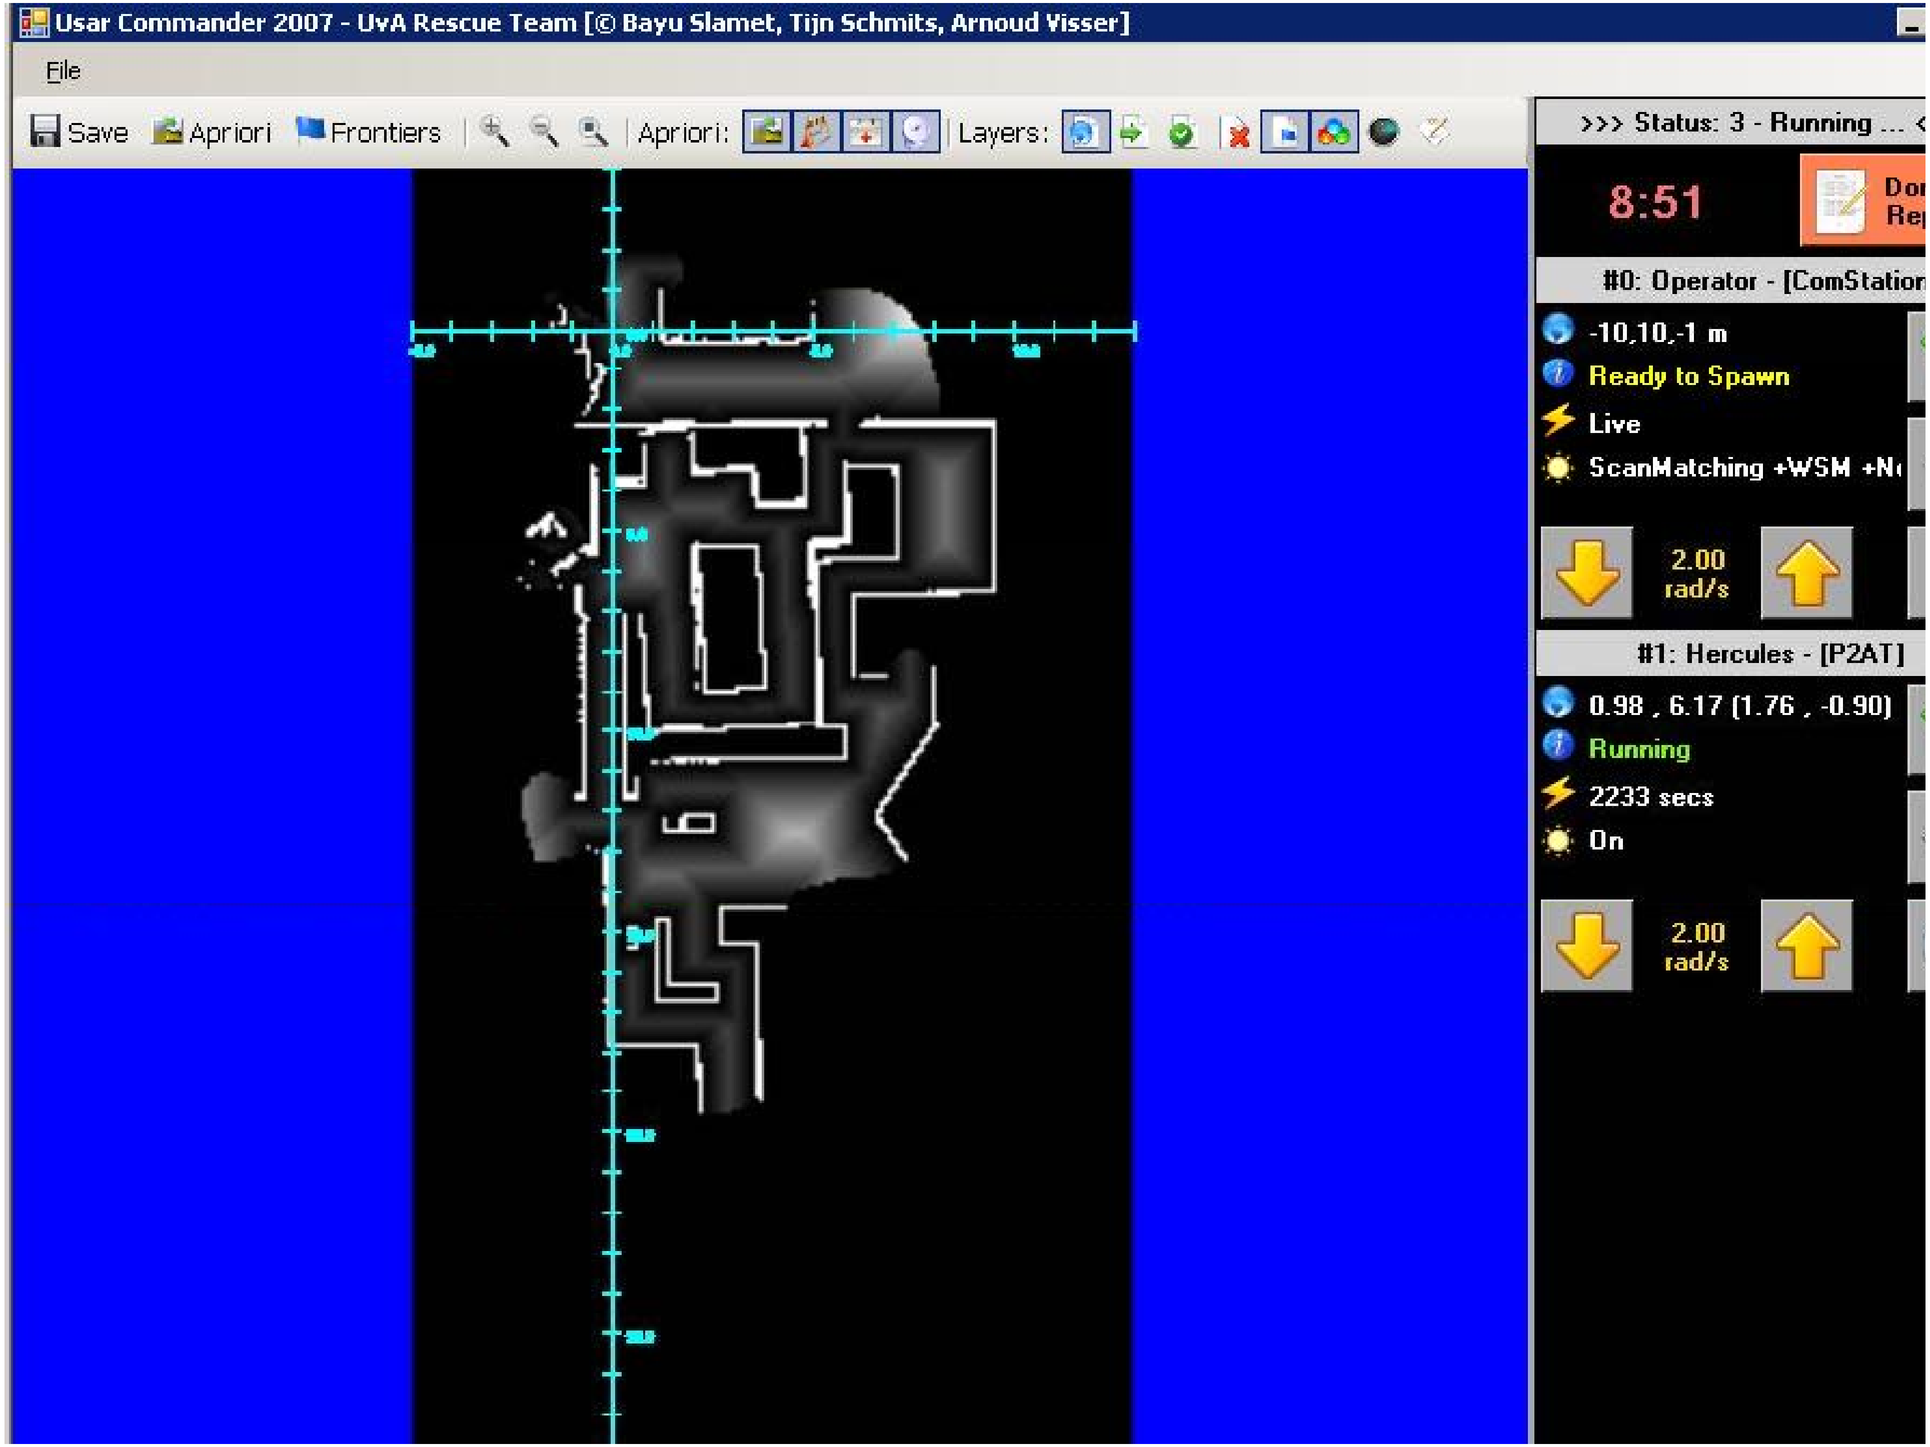

Usar Commander 2007 - UvA Rescue Team [© Bayu Slamet, Tijn Schmits, Arnoud Visser]
File
Save
Apriori
Frontiers
Apriori:
Layers:
>>> Status: 3 - Running ...
8:51
#0: Operator - [ComStation
-10,10,-1 m
Ready to Spawn
Live
ScanMatching +WSM +N
2.00
rad/s
#1: Hercules - [P2AT]
0.98 , 6.17 (1.76 , -0.90)
Running
2233 secs
On
2.00
rad/s

Usar Commander 2007 - UvA Rescue Team [© Bayu Slamet, Tijn Schmits, Arnoud Visser]
File
Save
Apriori
Frontiers
Apriori:
Layers:
>>> Status: 3 - Running ...
15:11
Frontiers
#0: Operator - [ComStation
-10,10,-1 m
Ready to Spawn
Live
ScanMatching +WSM +N
2.00 rad/s
#1: Hercules - [P2AT]
6.04 , 11.55 (6.91 , 4.18)
Paused
3297 secs
On
2.00 rad/s
30.96 m^2

Result for the 2006 Competition map

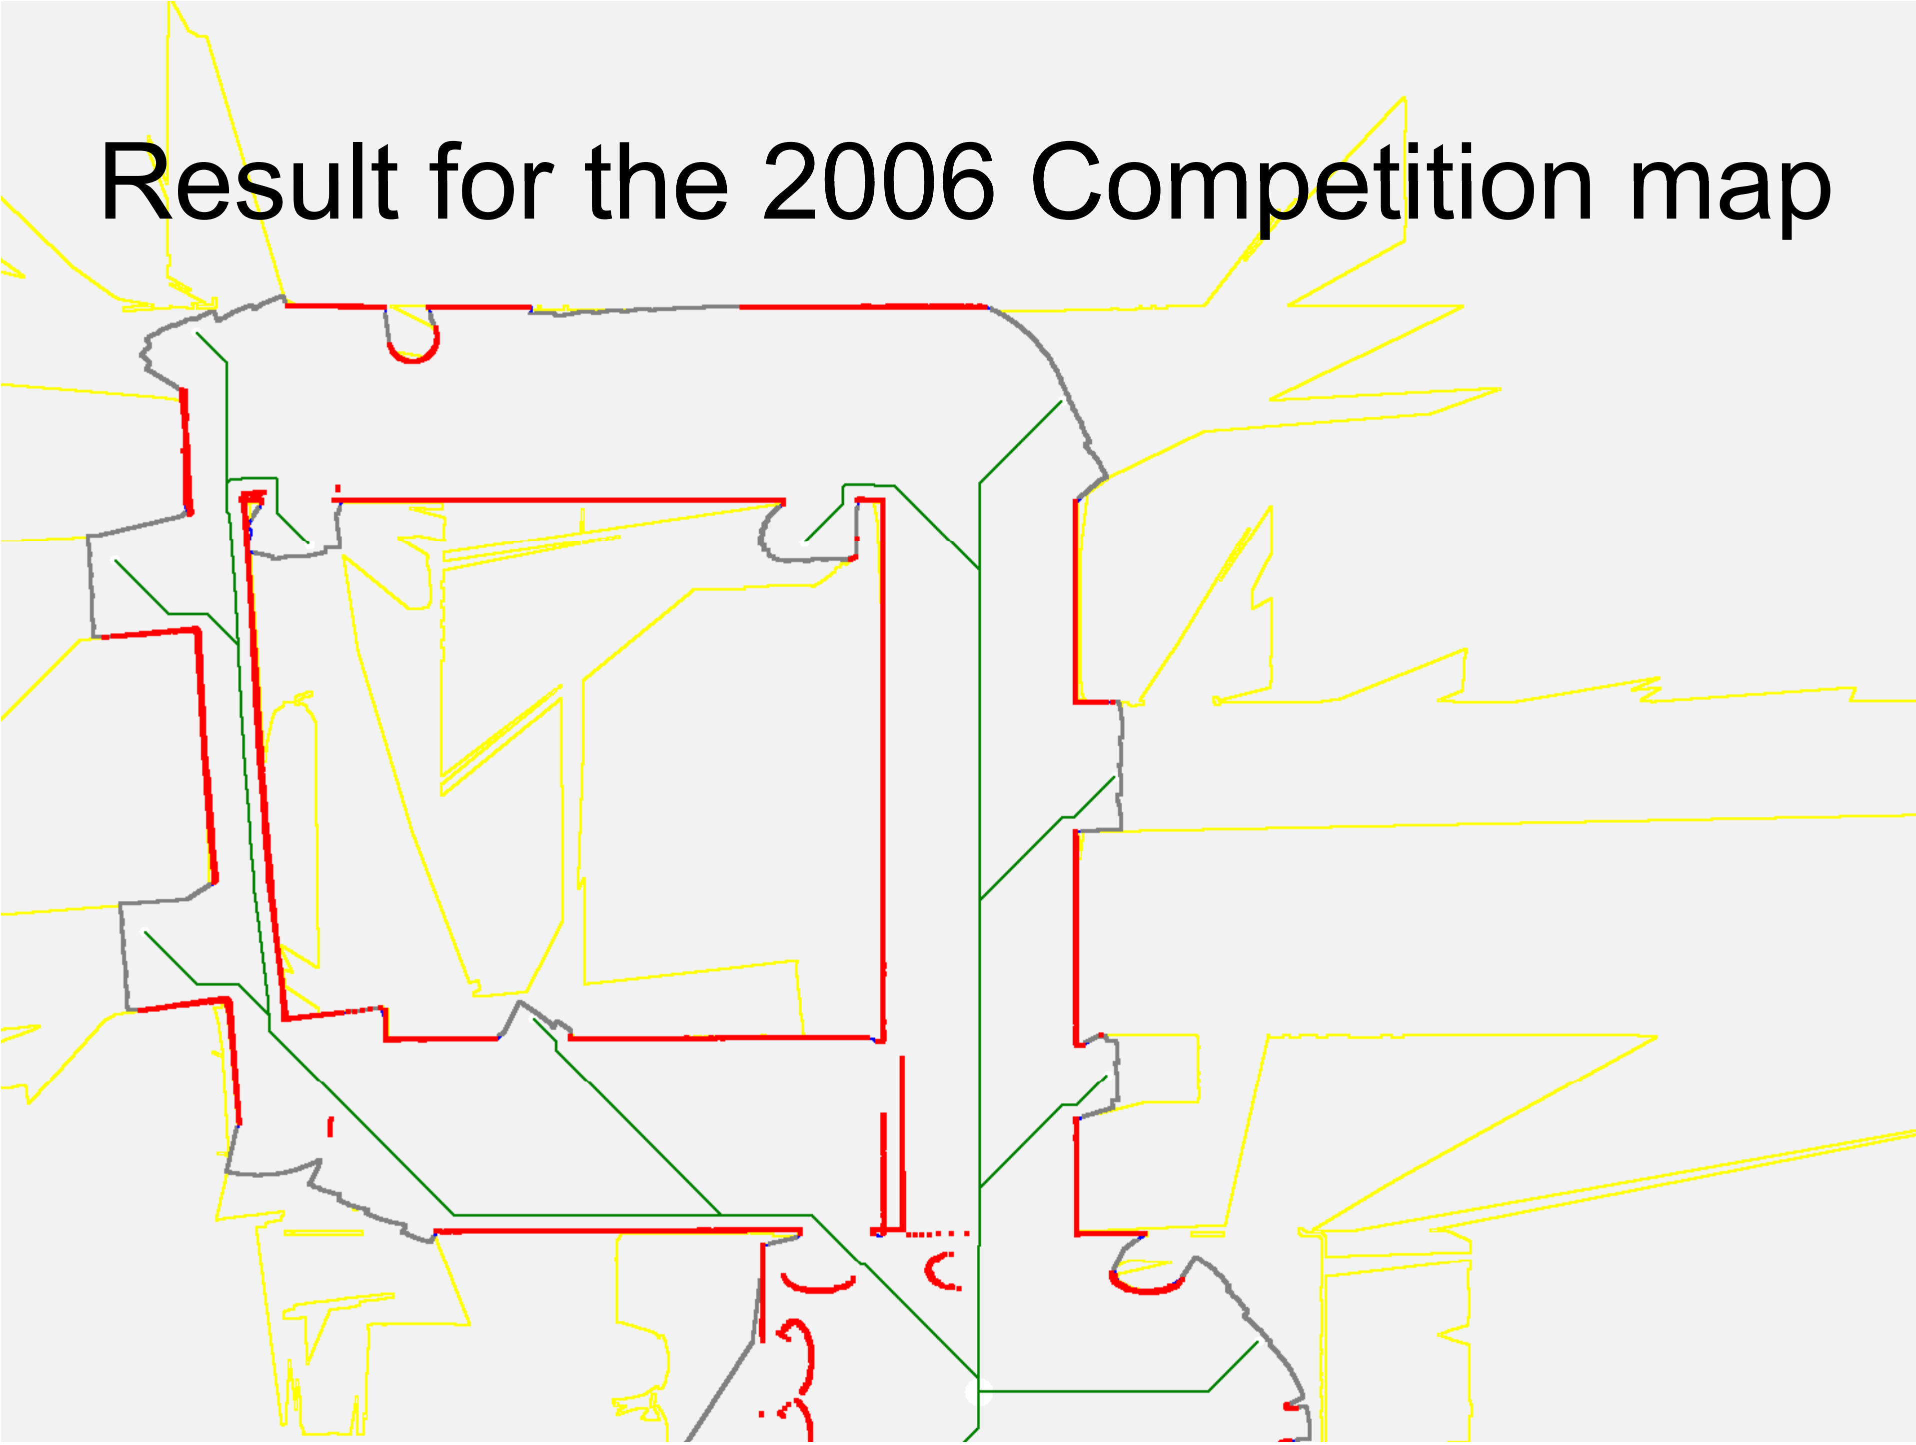

Quantative results

	expected frontiers & observation points	**exploration frontiers** found	*false positives*	*false negatives*	**observation points** found	*false positives*	*false negatives*
Three corridors (Fig. 2)	22	27	6	1	21	0	1
Lobby loop (Fig. 5)	13	17	5	1	11	0	2
Yellow arena (Fig. 6)	9	17	8	0	9	1	1

Navigation Challenge:
Result

- A limited set of candidate exploration locations can be generated from the current map *m*
- The candidate exploration locations are at the center of exploration frontiers
- The information gain can be estimated from the area beyond the frontiers

Balancing the information gain against the movement cost for multi-robot frontier exploration

Arnoud Visser and Bayu A. Slamet.

2nd European Robotics Symposium,
Prague, Czech Republic March 26, 2008

Universiteit van Amsterdam
Informatica Instituut
intelligent autonomous systems

Exploration Algorithm

- A greedy algorithm that calculates the utility function *U(f)* for each frontier *f* based on the *Area(f)* and distance *dist(f)*

$$U(f) = Area(f) / dist(f)$$

- The *Area(f)* is a measure for the information gain, and the *dist(f)* is a measure for the exploration cost.

Quantative results

	Single robot	Two robot team	Increase
North-East	416 m^2	519 m^2	+25%
North-West	379 m^2	637 m^2	+68%
South	498 m^2	707 m^2	+42%

Three Exploration Algorithms

- A greedy algorithm that uses an utility function *U(f)* for each frontier *f.* The utility function balances the information gain against the exploration costs.
  - linear: $U(f) = Area(f) / dist(f)$
  - quadratic: $U(f) = Area(f) / dist(f)^2$
  - cubic: $U(f) = Area(f) / dist(f)^3$
- The *Area(f)* is a measure for the information gain, and the *dist(f)* is a measure for the exploration cost.

Results

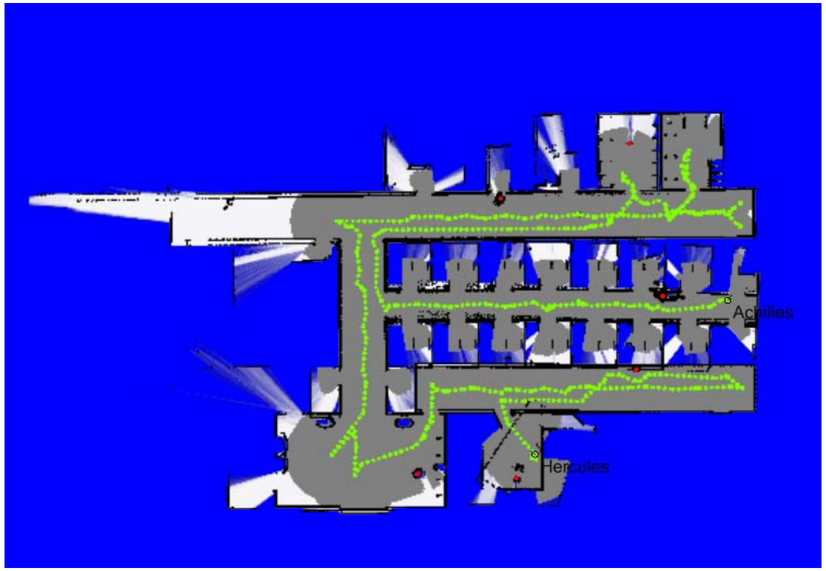

A(f)/d(f)
(629 m2, 6 victims)

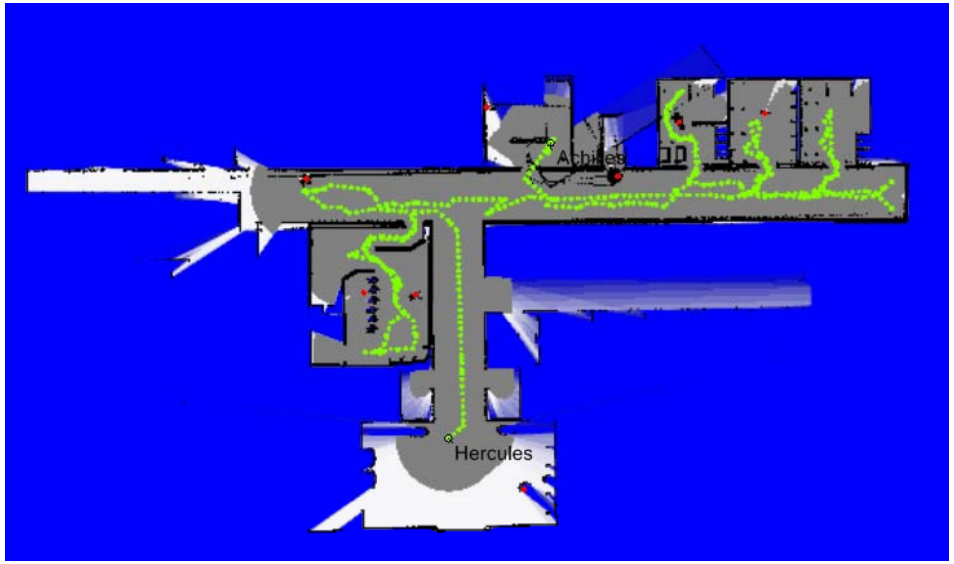

A(f)/d^3(f)
(486 m2, 8 victims)

Conclusion

- The information gain *A(f)* and exploration costs *d(f)* can in real-time be estimated from the current map
- Tuning the balance between the information gain and exploration costs can change the **overall behavior** from exploring mainly corridors towards concentrating on nearby rooms.

Including Communication Success in the Estimation of Information Gain for Multi-robot Exploration

Arnoud Visser and Bayu A. Slamet.

1th Workshop on Wireless Multihop Communications
in Networked Robotics,
Berlin, Germany, April 4, 2008

Universiteit van Amsterdam
Informatica Instituut
intelligent autonomous systems

Including the success of communication into exploration algorithm

- A greedy algorithm that uses an utility function *U(f)* for each frontier *f.* The utility function balances the information gain against the exploration costs.

$$U(f) = Area(f) \cdot P_{com} / dist(f)$$

- The *Area(f)* is a measure for the information gain, is now multiplied with the chance that gained information can be communicated to the ComStation.

Exploration Algorithm

distributed:
(map m as known by r_c)

double loop:
(first Euclidian distance,
followed by PathPlanning for u_{max})

Data: the identity of the current robot $r_c \in R$ and the map m as known by r_c.
Data: the set of robots r_i *in* R. Each r_i consist of the tuple $(x_{r_i}, y_{r_i}, \theta_{r_i})$.
Data: the set of frontiers f_j *in* F. Each f_j consist of the tuple $(x_{f_j}, y_{f_j}, A_{f_j}, P_{f_j})$.
Result: the pair r_c, f_c and the path p_c to the location (x_{f_c}, y_{f_c})

for *each robot* r_i *in* R **do**
  for *each frontier* f_j *in* F **do**
    $d_{eu} - \sqrt{(x_{f_j} - x_{r_i})^2 + (y_{f_j} - y_{r_i})^2}$;
    $u_{ij} = A_{f_j} P_{f_j} / d_{eu}$;
  end
end
$u_{max} = \max u_{ij}$;
repeat
  for *robot* r_i *and frontier* f_j *of* u_{max} **do**
    p=PathPlanning from (x_{r_i}, y_{r_i}) to (x_{f_j}, y_{f_j}) on map m;
    d_{pp}=length of path p;
    $u_{ij} = A_{f_j} P_{f_j} / d_{pp}$;
  end
  if $\max u_{ij} = u_{max}$ **then**
    Assign f_j to r_i;
    Prune U from i and j;
  end
  $u_{max} = \max u_{ij}$;
until *robot* r_c *is Assigned* ;
p_c=last path p

Estimating the Success of Communication

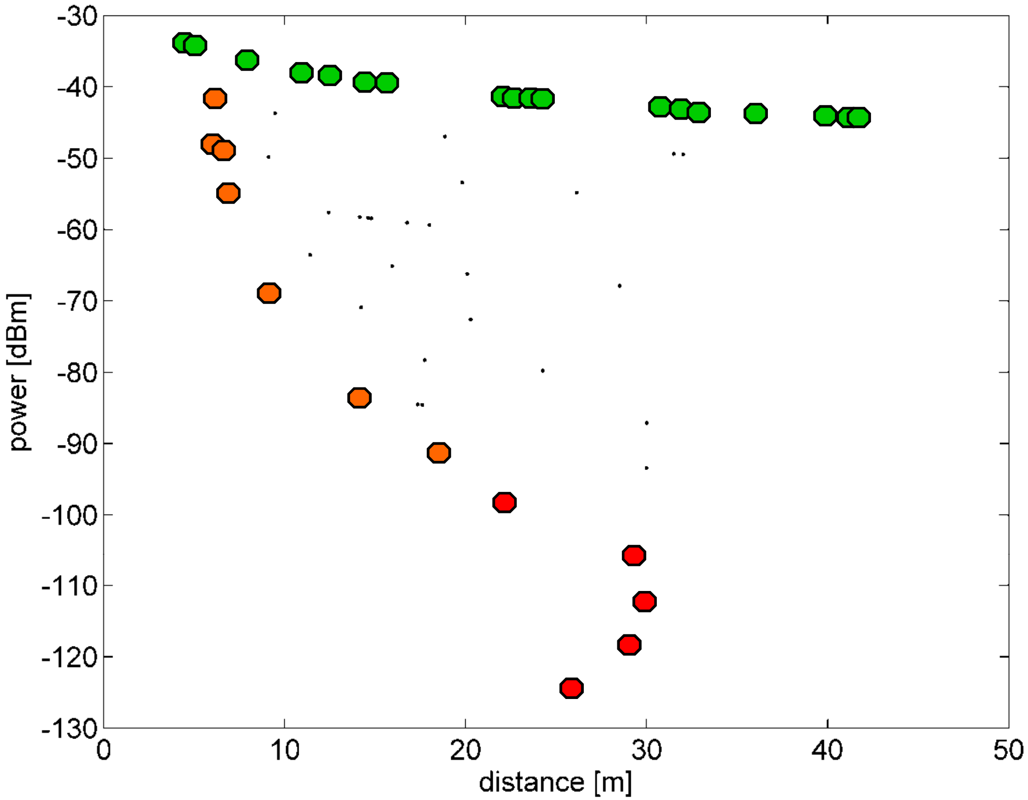

- The relation between distance and power is learned while exploring (depended on attenuation of walls)

Results

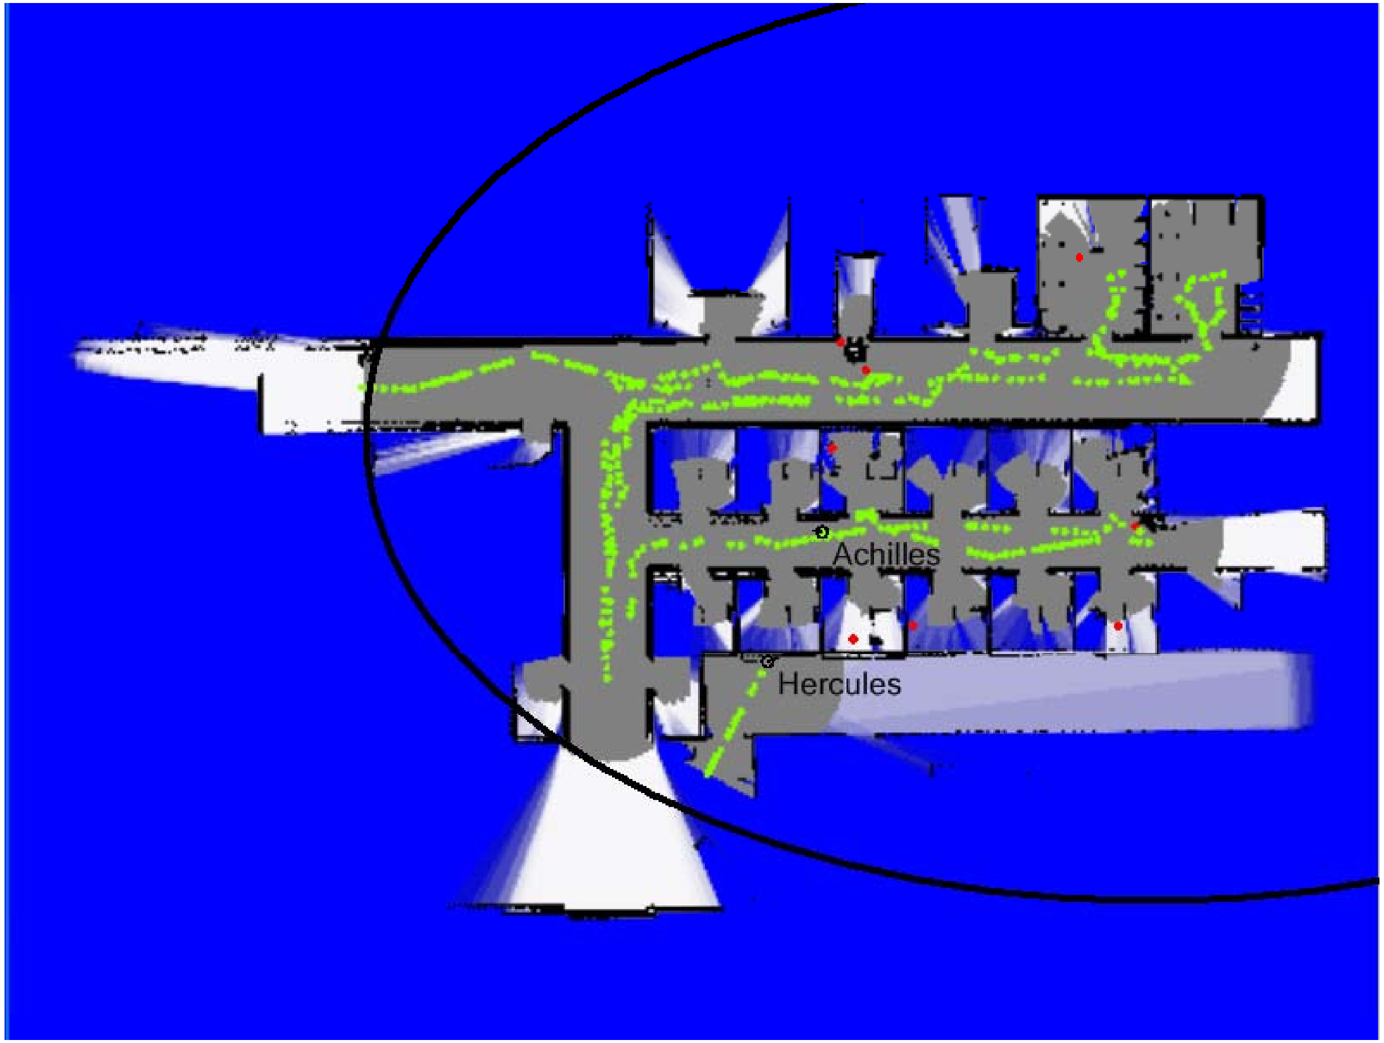

Two robots with ComStation in North-East corner,
robots do <u>not</u> share their map (545 m2, 6 victims)

Results

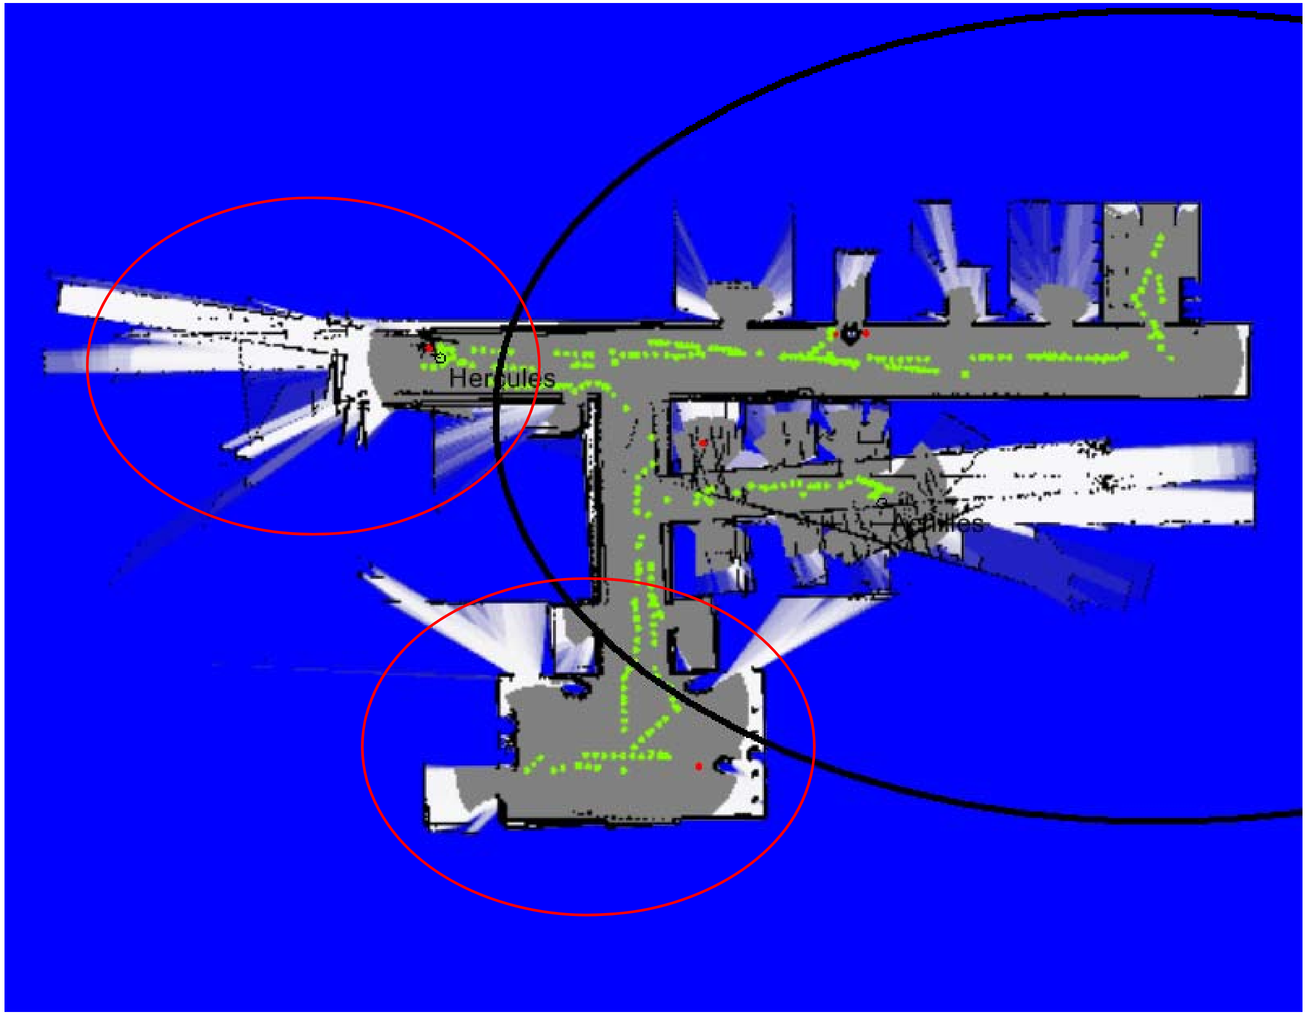

Two robots with ComStation in North-East corner, robots <u>share</u> their map (568 m2, 4 victims)

ComStation T-junction

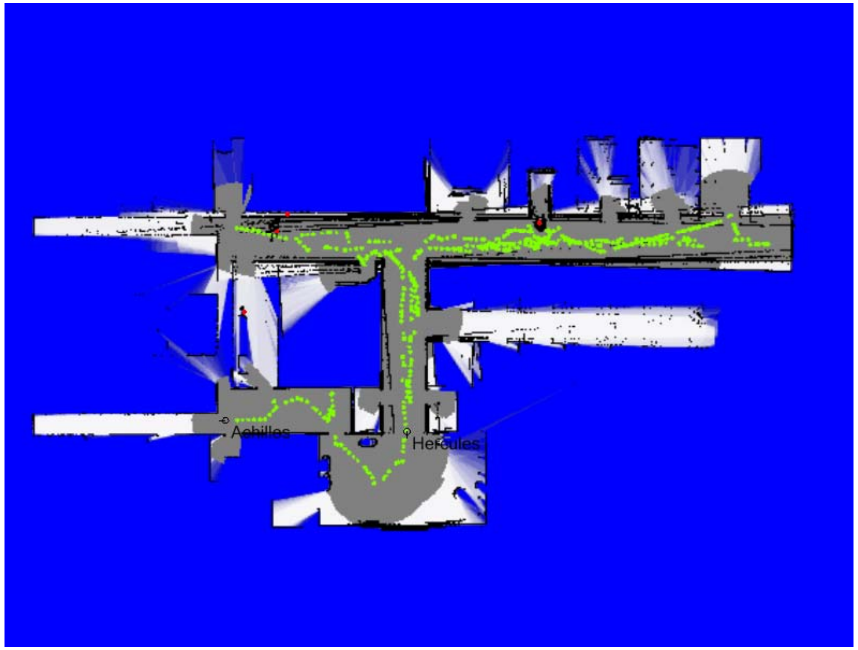

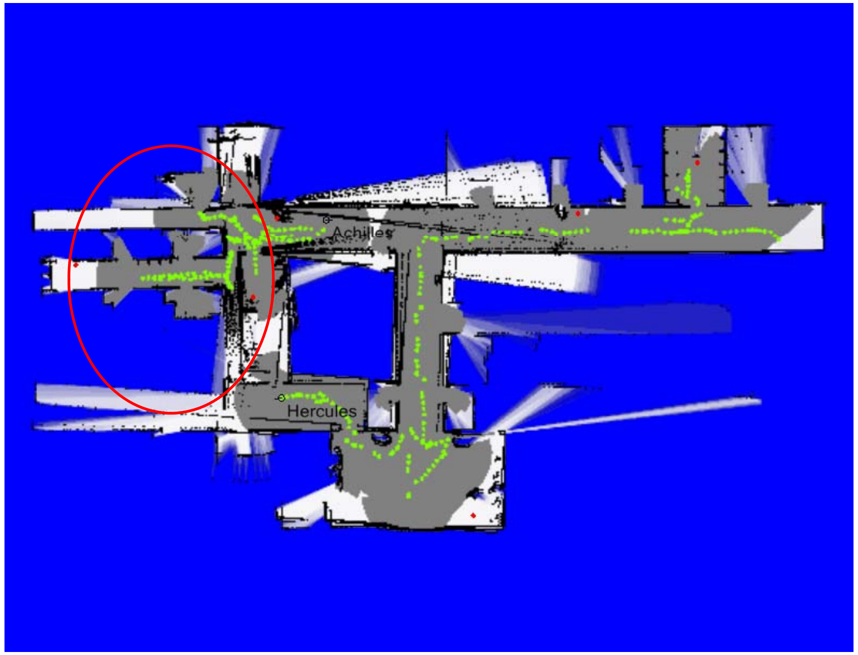

<u>not</u> share their map (556 m2, 4 victims), <u>shared</u> (642 m2, 5 victims)

ComStation NW

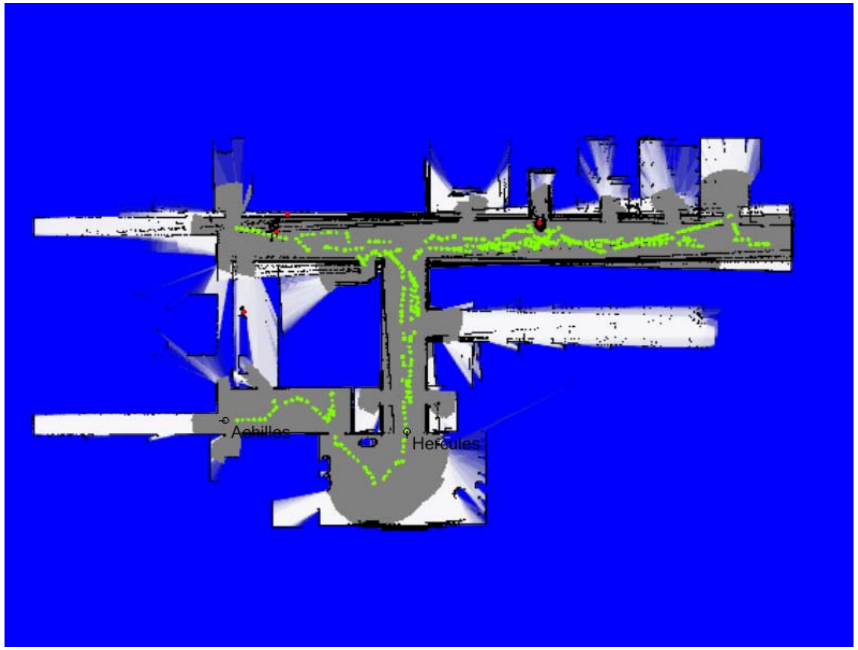

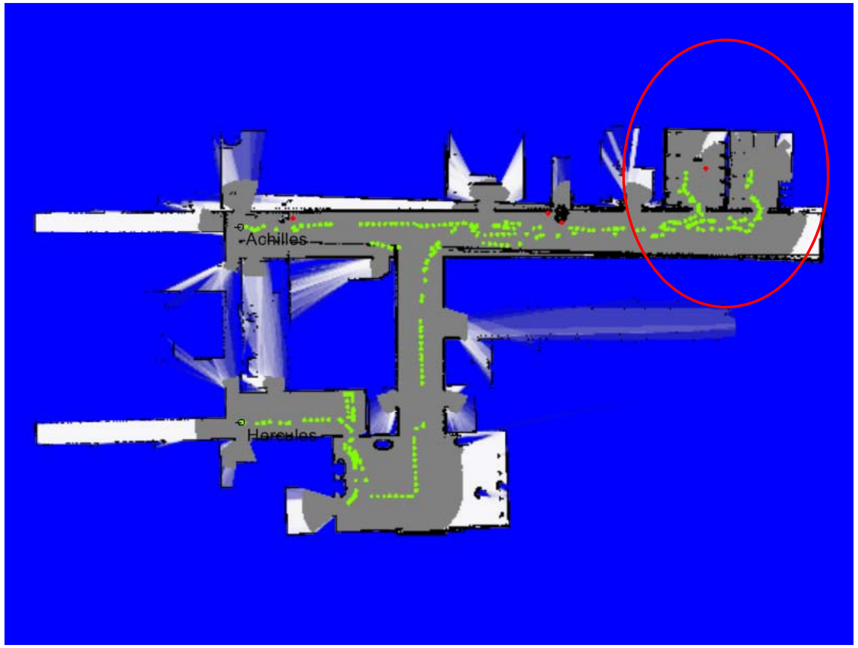

<u>not</u> share their map (535 m2, 3 victims), <u>shared</u> (531 m2, 3 victims)

Lesson learnt

- With shared maps less area is explored multiple times
- More information is collected in the outer regions, outside the direct communication range.
- More exploration, also means more navigation hazards

Improved navigation

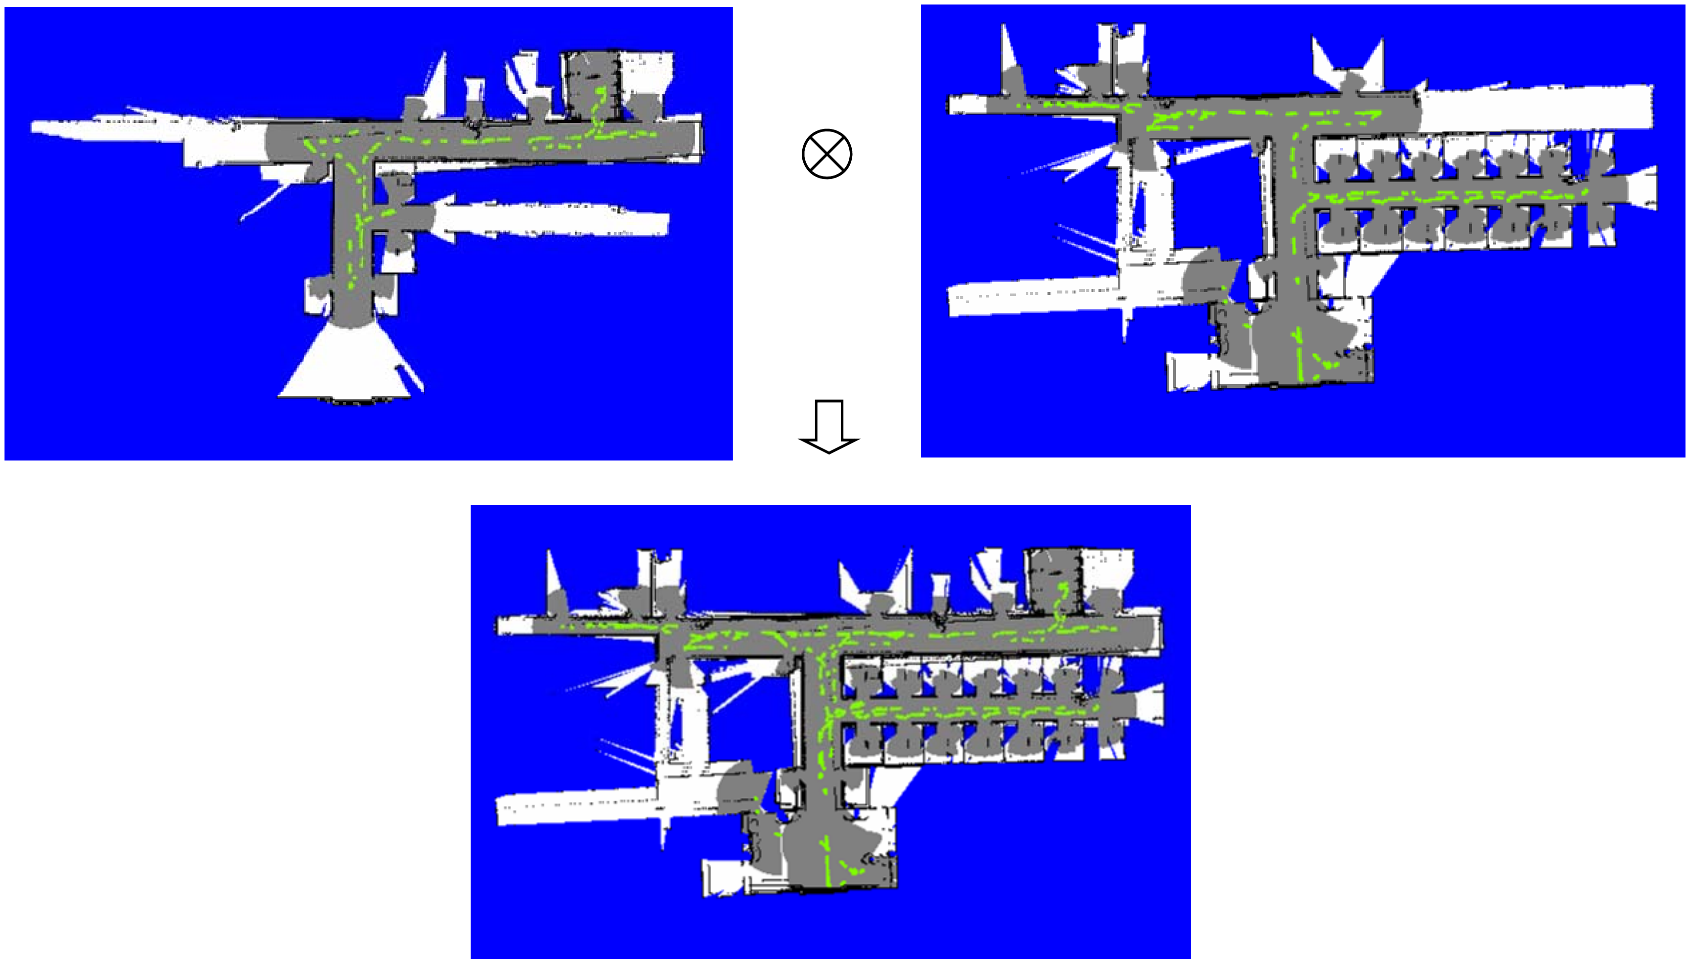

Shared map (879 m2, 5 victims)

Conclusion

- An exploration algorithm is introduced that includes the probability of communication success in the utility-function
- In a number of experiments, the need to include this probability inside the planning of multi-robot exploration is demonstrated.

Discussion

- Planning far ahead makes it difficult to estimate the probability of communication success at that distance
- Planning less far ahead could be beneficial for the communication, but not for the exploration
- We had hoped that exploration outside the communication range would occur for both robots at the same time, and that both robots could benefit from each other efforts.
- The limited number of experiments indicates that our hope about the timing could be correct.
- Yet, typically the robots are not near at that moment, so there is no guarantee that both robots could benefit from each other.

Planned Cooperation

Assistant communication relay roles have to be planned:

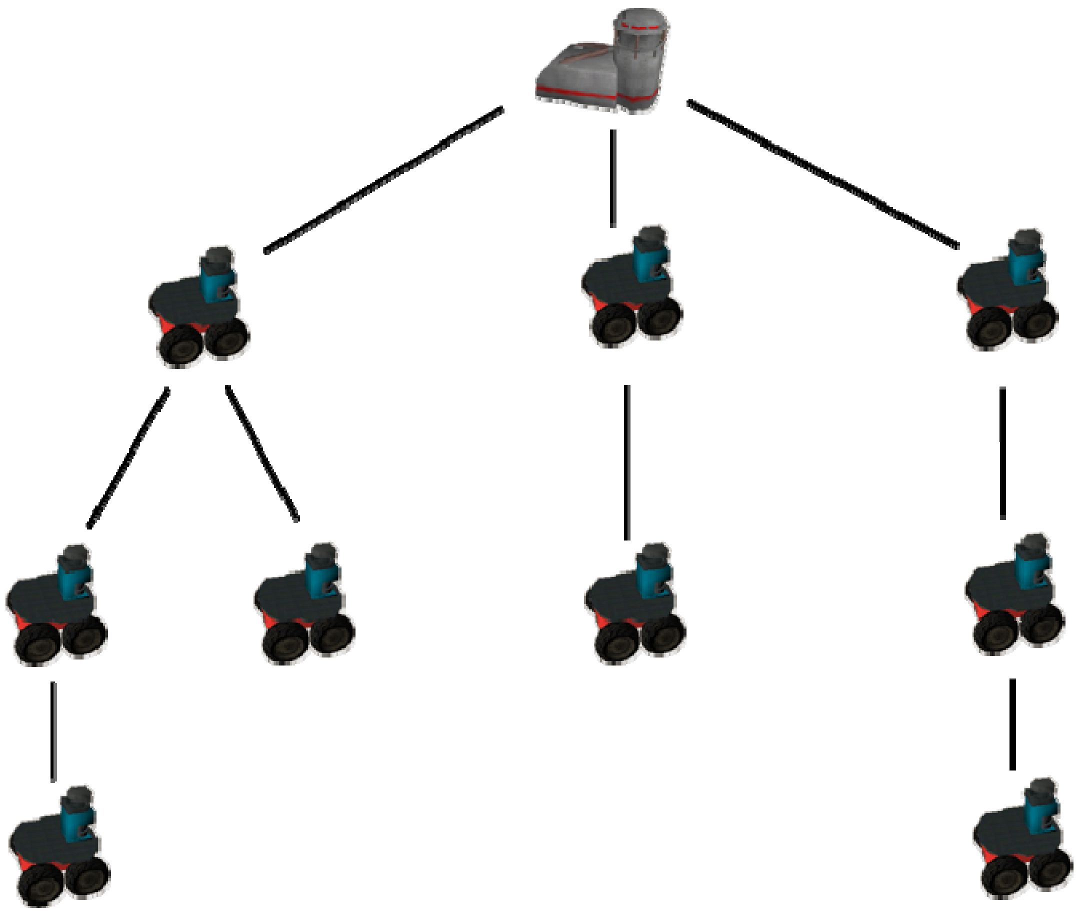

Conclusion

- An exploration algorithm is introduced that includes the probability of communication success in the utility-function
- In a number of experiments, the need to include this probability inside the planning of multi-robot exploration is demonstrated.

ROBOCUP ATLANTA 2007

Tele-operation

RoboCup Brazil Open 2008

Fully autonomous