

# Automated Task-Based Synthesis and Optimization of Field Robots

Chris Leger

The Robotics Institute  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
email: blah@cmu.edu  
phone: (412) 268-8157  
fax: (412) 268-5895

John Bares

The Robotics Institute  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
email: bares@cs.cmu.edu  
phone: (412) 268-7091  
fax: (412) 681-6961

## Abstract

*We present Darwin2K, a widely-applicable, extensible software tool for synthesizing and optimizing robot configurations. The system uses an evolutionary optimization algorithm that is independent of task, metrics, and type of robot, enabling the system to address a wide range of synthesis problems. Darwin2K can synthesize fixed-base and mobile robots (including free-flying robots, mobile manipulators, modular robots, and multiple or bifurcated manipulators), and includes a toolkit of simulation and analysis algorithms which are useful for many synthesis tasks. Some of these capabilities, such as dynamic simulation, are novel in automated synthesis of robots. An extensible software architecture enables new synthesis tasks to be addressed while maximizing use of existing system capabilities; this extensibility is a key contribution of the system. A key challenge is effectively optimizing multiple performance metrics; we present a method called Requirement Prioritization that guides the evolutionary algorithm through the design space. We apply Darwin2K to a robot synthesis tasks that includes synthesis of robot kinematics, dynamics, structural geometry, and actuator selection to meet and optimize multiple performance requirements.*

## Introduction: Why Automated Synthesis?

Robot configuration design is characterized by generating an artifact which is capable of motion in an area or volume. Typically, the configuration process generates the overall form of the robot, including kinematics and other geometry at the bare minimum but often including approximate descriptions of inertial properties, actuator and material selection, and structural geometry.

This design process is often performed in an ad hoc manner. It can be difficult to translate the requirements of a task into a robot configuration, and in contrast to some other engineering disciplines there are few design rules that can simplify this process. Human designers rely on intuition and experience with related design problems, and on engineering rules based on the experience of other expert designers. When designing a robot for a task with many new characteristics, relevant experience may be quite limited and may unnecessarily restrict the range of designs that are explored.

Frequently, a human investigates a small number of concepts on paper and selects a few that look promising. More detailed studies may then be performed on these, culminating in the simulation of one or more designs. One of the candidates is selected for detailed design, with tools such as finite element analysis used to evaluate parts of the design. Once the robot is built, changes may be required due to unforeseen problems or interactions that were not modeled in simulation. Significant design iterations are often not practical, since much of a project's schedule and resources may be devoted to creating a single robot; building a second or third robot to remedy design flaws is out of the question for many robot design projects. Thus, it is crucial to perform as much analysis and simulation as possible before the robot is built, and it is highly desirable to "get it right" the first time -- since the first time may be the only time.

Because of these factors, automated synthesis tools are especially attractive for robot design. Synthesis tools can address design problems for which there are little or no relevant human experience, can explore much larger numbers and ranges of designs than a human, can quickly perform design iterations in simulation, and can produce a well-optimized solution with high confidence of performance, all of which contribute to the likelihood of success of the first physical implementation.

Darwin2K is a software toolkit for automated synthesis of robot configurations. It includes capabilities for quickly describing and modifying robot configurations, simulating configurations as they perform tasks, and automatically synthesizing configurations to meet task-specific requirements and to optimize performance. Darwin2K is very useful in the early stages of the configuration process, as it can automatically explore tens of thousands of designs and can allow a human designer to rapidly perform design iterations and evaluate potential robots in simulation.

## Related Work

There has been much prior work in the area of robot design. However, research in *automated* design--that is, the development of software systems which perform a significant part of the synthesis of a robot--

has been limited to a handful of systems with widely varying scope and goals.

Most configuration synthesis systems have employed evolutionary algorithms of some sort, due to the robustness of such algorithms to local minimal and to search spaces that are highly nonlinear and of varying dimension. Previous approaches to robot configuration synthesis can be divided into two categories: modular design and non-modular design. The former group is characterized by constructing robots from fixed modules, mirroring a set of reconfigurable hardware modules from which the actual robots are built. The latter group has synthesized robots that are not built from fixed modules, but which consist of purely kinematic descriptions (i.e. Denavit-Hartenburg parameters or equivalent).

Much of the previous work in synthesizing robot configurations has been in the area of modular robots. Most of these have been for modular manipulators ([Ambrose94], [Chen95], [Paredis96], [Chocron97], [Han97]), while one ([Farritor96]) addressed mobile robots. Several of these approaches stand out for their unique contributions: [Ambrose94], while limited to planar manipulators, included modeling of actuators and link deflection. [Paredis96] used higher-fidelity simulation (including collision detection for non-trivial geometry) and generated fault-tolerant manipulators and motion plans. [Farritor96] generated modular field robot configurations and motion plans, though a low-fidelity evaluation process was used for configuration synthesis in the interest of decreasing computation time.

Other work has addressed kinematic configuration of monolithic (non-modular) manipulators. Kim and Khosla [Kim93] use a genetic algorithm to determine the base location and Denavit-Hartenberg parameters for a manipulator. Chedmail and Ramstein [Chedmail96] use a genetic algorithm to determine the type (one of several manipulators) and location of a robot to optimize workspace reachability. McCrea [McCrea97] discusses the application of genetic algorithms to the selection of several parameters for a manipulator used in bridge restoration.

Several key limitations are apparent in previous robot synthesis systems. The only systems to consider non-kinematic properties were those using fixed modules; the systems for monolithic manipulators are purely kinematic and have limited robot representations--specifically, monolithic manipulators have been modeled as joints connected by zero-thickness line segments. The strictly modular systems, while able to represent non-kinematic properties, are not able to independently vary properties such as actuators or link structure, thus leading to suboptimal solutions. Previous systems have not addressed analysis and simulation needs such as dynamic simulation and estimation of link deformation due to loads, thus inhibiting the ability of these systems to meaningfully optimize actu-

ator selection, link structural geometry, and inertial properties.


## System Description

An earlier version of Darwin2K is described in [Leger98]. Here, we will briefly discuss Darwin 2K's architecture and describe recent additions, including an improved method for optimizing multiple metrics, dynamic simulation, and optimization of non-kinematic properties.

Darwin2K consists of two distinct parts: a configuration synthesizer, and a simulation and analysis tool. The synthesizer, called the Evolutionary Synthesis Engine (ESE), uses the simulation tool to evaluate the performance of each new configuration it creates. The ESE's evolutionary algorithm synthesizes robot configurations by applying genetic operators to one or more existing parent configurations. The ESE selects the parents from a population of configurations according to their performance, so that each new configuration is likely to perform well and will occasionally outperform its parent configuration(s). The performance of each robot is measured in a task-specific manner: the robot performs the task in simulation, and multiple task-specific metrics (selected by the designer) record various aspects for the robot's performance. Darwin2K's 15 existing metrics include power consumption, task completion time, robot mass, stability, actuator saturation, and collision measurement.

Each configuration is assembled from one or more *parameterized modules*, of which there are four basic types: bases, links, joints, and tools. Each module represents a part of a robot, with an arbitrary number of parameters that dictate specific properties. Modules can have varying complexity, ranging from a simple link with no moving parts, to a joint with one or degrees of freedom, to an entire manipulator, to a mobile base. Each module can have a number of connectors, which indicate how the module can be attached to other modules. A module's parameters may describe any property of the module, such as a geometric dimension, a discrete component selection (e.g. a motor, gearhead, or material choice), or a controller gain. All modules have the same generic interface so that new module types can be added without requiring changes to the ESE, thus enabling Darwin2K to use (if necessary) task-specific modules when addressing design problems. Darwin2K currently contains approximately 30 general-purpose parameterized modules, as well as 10 or so task-specific modules that were created for specific synthesis problems. See Figure 1 for examples of Darwin2K's parameterized modules.

To represent a particular robot configuration, one or more parameterized modules are assembled into a graph, called the Parameterize Module Configuration Graph (PMCG). Each node in the graph is a parameterized module, and each edge is a connection between modules. For serial-chain manipulators, this graph is



**Figure 1:** A selection of Darwin2K modules. From left to right:

- a SCARA module with 3 degrees of freedom and with parameters for motors, gearboxes, and joint angle offsets for each of three joints, and parameters for material, link lengths, link diameters, and wall thicknesses for the two links.
- a mobile base for a construction robot, with parameters for wheelbase, tread, and attachment location.
- a joint module that included parameters for motor, gearbox, and material selections, and for link length, tube diameter, and wall thickness.

simply a connected series of modules, usually starting with a base, followed by one or more links or joints, and ending with a tool. Multiple and branching manipulators can also be represented by having multiple branches in the graph. One important aspect of the PMCG is that each connection, and each module parameter, have a *const-flag*, which can be set to indicate to the ESE that a particular feature should not change. Thus, if the designer knows of any task-specific features that would be beneficial, they can be directly encoded and preserved by the synthesizer.

The PMCG representation is easy for a human designer to manually modify, thus allowing the designer to quickly build configurations that can be simulated using Darwin2K's simulation tool. In addition to displaying simulations of robots performing tasks, the simulator also operates in a no-display mode, so that the synthesizer can use it to measure the performance of configurations. Many simulators can be distributed over a network of computers, all providing fitness measurements to the synthesizer.

### Task Specification

Darwin2K performs task-specific synthesis: robots are synthesized for a specific task, and are evaluated with respect to the performance metrics and requirements of the task. There are two components to the task specification: a description of how the task is performed, and a set of performance requirements or metrics.

Darwin2K contains a generic simulator for one or more manipulators (possibly mounted on a mobile base) following a series of endpoint trajectories. The simulator has a wide range of capabilities including kinematic and dynamic simulation, collision detection (using the RAPID library [Gottschalk96]), motion planning, and estimation of link deflections. While this simulator is suitable for a range of typical manipulation tasks, each target application may have unique

simulation and evaluation requirements. For this reason, we have structured Darwin2K's software architecture so that a new, task-specific simulator can be quickly constructed using Darwin2K's existing simulation components, which include collision detection, several controllers, payload models, a motion planner for planar mobile robots, and several trajectory representations. Task-specific simulators usually require little coding, normally just a main simulation loop: the evaluation components handle the numerical details, while the main simulator just controls which simulation capabilities and controllers are used over the course of the simulation.

New simulation, control, and evaluation capabilities can also be added, in the form of C++ code contained in a dynamic library. Darwin2K's object-oriented software architecture allows these new capabilities to interface with existing software components: for example, a new trajectory representation can be used with an existing controller. Once an appropriate simulator is constructed (including any new simulation or control components), the designer can use the synthesizer to begin creating a robot for the task at hand; no modifications to the synthesis engine are required.

The vast majority of design problems involve multiple, often-conflicting metrics such as cost versus performance, speed versus power, or mass versus capability. Additionally, most tasks have hard limits that must be met for one or more metrics: there can be no collisions during operation, mass must be under a certain limit, etc. These requirements and non-linear dependencies (i.e. power consumption is irrelevant if a robot can't complete the task) cannot be effectively encoded by simple scalarization, such as using a weighted sum of metrics to create a single measure of performance. To remedy this, Darwin2K's task specification includes these constraints and priorities in a natural manner, called Requirement Prioritization. The designer selects a number of metrics (typically between 3 and 10), ranks them in order of importance, and sets an acceptability threshold for some. For example, a free-flying space robot tasked with servicing a satellite might have the following metric specification:

- 100% task completion
- 0 collisions
- < 3mm cross-track error at the tool endpoint
- average actuator torque < 80% of continuous torque rating for each actuator
- link deflection < 1mm
- minimal mass, energy, and task completion time.

Task completion, collision, and accuracy are most important: any configuration that does not satisfy these is inferior to any configuration that does. Thus, these are given a priority of 0 (most important). To ensure that actuators and links are appropriately sized, the actuator torque and link deflections are added next, with a

priority of 1. Finally, we would like to minimize mass, energy consumption, and completion time given that the other requirements are satisfied, so these are given a priority of 2. Darwin2K uses this specification to determine how configurations are selected for reproduction: the highest-priority metrics (completion, collision, and accuracy) are used until the population contains a certain number of configurations that meet the acceptability thresholds; then the next-highest group is used until enough configurations meet that group's thresholds; and finally the last group of metrics (mass, energy and time) are used until the synthesis process ends (since no acceptance criteria were set for them). The synthesis process halts when either a time limit is reached or when a fixed number of configurations have been generated (typically between 40,000 and 160,000). While optimizing within each group of metrics, the metrics are randomly selected for use based on how many configurations in the population satisfy each metric. Thus, the metrics whose acceptability thresholds are satisfied by few or no configurations will be used to select configurations more frequently than the metrics which many configurations satisfy. This automatically guides the synthesizer to remedy the weaknesses of the population as a whole. In practice, we have found this method to be much more effective in optimizing multiple metrics than scalarization approaches.

### Dynamic Simulation

While a purely kinematic simulation may be adequate for tasks that do not require large forces to be applied by or to the robot, there are some tasks for which dynamic simulation must be used to accurately assess a robot's performance. Two broad classes of robots for which dynamic simulation is crucial are free-flying robots (either in space or underwater), and large, high-force robots such as those targeted at construction applications. Free-flying robots experience reaction forces during manipulation that may move the entire robot, while construction robots frequently require one or more actuators to be operated at saturation (i.e. providing maximum force or torque output), in which case the response of the robot will not be known without a dynamic model.

Darwin2K supports forward and inverse dynamic modeling for fixed-base, mobile, and free-flying manipulators, including robots with branching or multiple manipulators. The inverse dynamic problem (computing joint torques required to produce desired joint accelerations) is relatively easy to solve using the well-known iterative Newton-Euler dynamic formulation [Craig89]. However, the forward dynamic problem--computing the motion of a robot given the force applied at each degree of freedom--is more difficult. The dynamic model of a robot can be expressed by the equation

$$T = M(\Theta)\ddot{\Theta} + V(\Theta, \dot{\Theta})$$

Where  $T$  is a vector of joint torques,  $M(\Theta)$  is the mass

matrix,  $V(\Theta, \dot{\Theta})$  is a vector of inertial and applied forces, and  $\Theta$  is a vector of joint positions. Forward dynamic simulation requires that  $M$  and  $V$  be constructed from the current configuration of the robot, after which  $\ddot{\Theta}$  is found by solving the set of equations. Computing  $M$  and  $V$  is the crux of the problem: the Newton-Euler equations (or equivalent) must be applied symbolically to yield equations for each element of  $T$ , and then these equations must be factored to pull out the coefficients for each entry of  $\ddot{\Theta}$ . Terms in the equations that do not contain a joint acceleration get lumped into  $V$ . While this can be done using Mathematica or a similar package, or even by hand for a known robot, these approaches are not very useful for automated synthesis where efficient dynamic simulation must be performed for the tens to hundreds of thousands of unique robots created by the synthesizer during a synthesis run.


Our approach to this problem is based on treating each scalar in the dynamic equations as a linear combination of joint accelerations and a constant term. We call these *s-val*s, for Separated Values, since each joint acceleration coefficient is accumulated separately. This is possible because the dynamic equations do not contain products of multiple  $\ddot{\Theta}$ 's. For example, a scalar *s-val* for a robot with  $n$  degrees of freedom would be a vector of length  $n+1$ :  $s = [m_0 \ m_1 \ \dots \ m_n]$ . Numerically evaluating  $s$  given a set of joint accelerations is performed by taking the dot product of  $s$  with an augmented vector of joint accelerations:

$$\begin{aligned} \text{value}(s) &= m_0 + m_1\ddot{\theta}_1 + m_2\ddot{\theta}_2 + \dots + m_n\ddot{\theta}_n \\ &= [m_0 \ m_1 \ m_2 \ \dots \ m_n] [1 \ \ddot{\theta}_1 \ \ddot{\theta}_2 \ \dots \ \ddot{\theta}_n]^T \end{aligned}$$

A 3D vector would thus be represented as 3 *s-val*s, one for each dimension. Vector operators for *s-val*s can be derived by symbolically applying normal vector operators (addition, cross-product, etc.) to a vector of *s-val*s and then re-grouping the terms for each  $\ddot{\theta}_i$  in the result. Note that a joint acceleration  $\ddot{\theta}_i$  is represented by an *s-val* with  $m_i = 1$  and all elements equal to zero (since multiplying this *s-val* with the augmented acceleration vector will yield  $\ddot{\theta}_i$ ).

To fill the  $i$ th row of  $M$  and  $V$ , we symbolically compute the equation for the  $i$ th joint torque, evaluate the equation numerically, and then extract the coefficient of each joint acceleration and the constant coefficient. Performing this process for each torque equation (i.e. for each row of  $M$  and  $V$ ) gives a system of equations that can be solved for  $\ddot{\Theta}$  given the torque or force applied at each joint ( $T$ ).

We have implemented a set of C++ arithmetic and vector classes specialized for this approach. Using these classes, the Newton-Euler equations are evaluated at each simulation time step. Each computed joint torque (i.e. element of  $T$ ) is an *s-val*; the *s-val*'s elements are the elements of the corresponding rows in  $M$  and  $V$ . After numerically evaluating the *s-val* for each joint torque equation (i.e. for each row of  $T$ ,  $M$ , and  $V$ ), we




**Figure 2:** (top to bottom) Best mass, energy, and task completion time of feasible configurations during optimization of final requirement group.

solve the system of equations for  $\ddot{\theta}$ . We then numerically integrate  $\dot{\theta}$  and  $\ddot{\theta}$  using the Runge-Kutta 4 algorithm [Press92] to compute the next state of the robot. We use an adaptive stepsizing algorithm to maximize numeric stability and computational efficiency.

### Design Examples

To verify Darwin2K's ability to perform configuration synthesis for challenging applications, we used Darwin2K to configure a free-flying space robot with two manipulators for a satellite servicing task. The task is loosely based on the requirements of the Ranger Telerobotic Flight Experiment and Telerobotic Shuttle Experiment ([SSL99]), and consists of both kinematic and dynamic simulations. The kinematic simulation is designed to ensure that the robot's manipulators have adequate workspaces, and consists of having the manipulators follow trajectories that cover a reasonable workspace. Collision detection is performed to ensure that there is no interference between the robot's manipulators.


The dynamic simulation is more involved and is designed to be representative of the operations required of the robot. The robot starts at a position near a 9100kg satellite, uses one of its two arms to grapple the satellite and then move the base to a known location relative to the satellite. The robot's other manipulator (the work manipulator) then removes an Orbit Replacable Unit from the satellite, and the grapple manipulator moves the base relative to the satellite again. Finally, the work manipulator re-inserts the ORU into the satellite at another location. Darwin2K's free-flyer controller (which is based on the Singularity-Robust Inverse [Nakamura86] and augmented with the robot's dynamic model) was used to control the robot during the dynamic simulation. The ORU has a mass of 5kg and requires a torque of 14 Nm (10 ft-lbs) to en-



**Figure 3:** Screenshot from Darwin2K's simulator showing the configuration with lowest energy consumption as it begins to insert the ORU into the satellite.

gage or disengage from the satellite, and 45N (10lbs) of force to insert or remove. The metrics used were those specified earlier in the "Task Specification" section. During optimization of the first requirement group (task completion, collision detection, and accuracy), kinematic simulation was used for the entire simulation; dynamic simulation was used as specified above for the remaining two requirement groups.

Five module types were used in the synthesis process: a free-flying base module with a mass of 267kg; a tool module representing Ranger's Microconical End Effector (MEE); and three joint modules (each with five parameters for various dimensions and for motor and gearbox selections). The robot was constrained to have two symmetric arms with 7 or 8 degrees of freedom each, and in addition to the parameters of each module, there were six task parameters: two specified the base's initial pose relative to the satellite, and the remaining four specified the velocity and acceleration of the base during each of the two relative base motions. The design space contained  $5.14 \times 10^{58}$  configurations. For this problem, we ran Darwin2K using idle CPU time from approximately 30 computers over 8 hours. The initial population of 200 configurations contained 190 unique PMCG topologies. After evaluating 5,000 configurations, 100 out of 200 configurations in the population satisfied the first requirement group. After 11,000 evaluations, a configuration satisfying the second requirement group had been created, and Darwin2K moved to the final requirement group at 20,000 configurations. At this point, 16 of the 200 configurations in the population satisfied the task requirements; the best of these all had the same PMCG topology. The remainder of the run was spent minimizing mass, energy, and task completion time, primarily through parametric variation: at the end of the run, the population contained only a few distinct topologies, with 98% of the configurations sharing the same topology. Thus, the final stage of optimization focused on parametric, rather than topological, variation. Figure 2 shows the best mass, energy, and time of the feasible configurations as the final requirement



**Figure 4:** A gantry-mounted SCARA manipulator synthesized for a storage and retrieval task (upper left), and a 5-DOF arm with prismatic joint synthesized for waterproofing tiles on the Space Shuttle (lower right).

group was optimized. During the last requirement group, the mass of the lightest feasible configuration in the population decreased from 541kg to 418kg; the lowest energy consumption decreased from 2.1kJ to 0.6kJ, and task completion time decreased from 41s to 37s. Figure 3 shows the feasible configuration with lowest energy consumption. After 8 hours, 43,000 configurations had been evaluated; on average, each evaluation took 20 seconds of real-time for roughly 40 seconds of simulated time.

We have also applied Darwin2K to other synthesis tasks where dynamic simulation is not required. In these cases, actuator selection and link structural sizing can still be performed since the inverse dynamic equations can be applied at each time step to estimate the torque required at each actuator. These torques can be compared to the torque capability of each actuator, and can be used to estimate link deflection to ensure that links are sufficiently stiff. Figure 4 shows results from two of these other synthesis tasks: a robot for storing and retrieving containers of wafers in semiconductor fabrication plant, and a manipulator for waterproofing Space Shuttle tiles (based on the task described in [Kim93]). While the details of these synthesis tasks are beyond the scope of this paper, they are mentioned here to give the reader an idea of the breadth of application of Darwin2K.

## Conclusion

Darwin2K is a practical, widely-applicable task-based automated synthesis system for robots, and is significantly more capable than previous approaches. The system's extensible, modular software architecture allows task-specific simulations to be easily constructed so that new tasks may be addressed in a relevant manner. Darwin2K's synthesis engine is independent of task and robot type, and can synthesize robots for complex or poorly-understood applications.

A key challenge for automated synthesis is generating designs which meet and optimize multiple performance requirements; Requirement Prioritization is a response to this need and has worked well for us in practice. Dynamic simulation is another important as-

pect of robot synthesis. Our method of dynamic simulation enables synthesis of robots for tasks where dynamic behavior is important, such as microgravity, high-force, or high-speed applications. Future work will concentrate on expanding Darwin2K's library of simulation components, applying it to real-world design problems, and packaging it for widespread use.

## References

- [Ambrose94] R. Ambrose and D. Tesar. The Optimal Selection of Robot Modules for Space Manipulators. In *Proceedings of the ASCE Space 94 Conference*.
- [Chedmail96] P. Chedmail and E. Ramstein. Robot mechanism synthesis and genetic algorithms. In *Proceedings of ICRA 1996*, pages 3466–3471.
- [Chen95] I. Chen and J. Burdick. Determining task optimal modular robot assembly configurations. In *Proceedings of ICRA 1995*, pages 132–137, 1995.
- [Chocron98] O. Chocron and P. Bidaud. Genetic Design of 3D Modular Manipulators. In *Proceedings of ICRA 1997*, pp. 223–228.
- [Craig89] J. Craig. *Introduction to Robotics: Mechanics and Control*. Addison-Wesley, 2nd edition, 1989.
- [Gottschalk96] S. Gottschalk, M. C. Lin and D. Manocha. *OBTree: A hierarchical Structure for Rapid Interference Detection*. Technical Report TR96-013, Department of Computer Science, University of North Carolina, Chapel Hill, 1996.
- [Han98] J. Han, W. K. Chung, Y. Youm, and S. H. Kim. Task Based Design of Modular Robot Manipulator using Efficient Genetic Algorithm. In *Proceedings of ICRA 1997*, pp. 507–512.
- [Kim93] J.-O. Kim and P. Khosla. Design of space shuttle servicing robot: An application of task-based kinematic design. In *Proceedings of ICRA 1993*, pp 867–874.
- [Leger98] C. Leger and J. Bares. Automated Synthesis and Optimization of Robot Configurations. In *Proceedings of the 1998 ASME Design Engineering Technical Conferences*.
- [McCrea97] A. McCrea. Genetic algorithm performance in parametric selection of bridge restoration robot. In *Proceedings of the 14th International Symposium on Automation and Robotics in Construction*, pages 437–441, 1997.
- [Nakamura86] Y. Nakamura. Inverse kinematic solutions with singularity robustness for robot manipulator control. *Journal of Dynamic Systems, Measurement, and Control*, 108:163–171 September 1986.
- [Paredis96] C. Paredis. *An Agent-Based Approach to the Design of Rapidly Deployable Fault Tolerant Manipulators*. PhD thesis, The Robotics Institute, Carnegie Mellon University, 1996.
- [Press92] W. Press and S. Teukolsky and W. Vetterling and B. Flannery, *Numerical Recipes In C*. Cambridge University Press, 2nd edition, 1992.
- [SSL99] University of Maryland Space Science Laboratory's Ranger project homepage: <http://www.ssl.umd.edu/homepage/Projects/ranger.html> (as of July 1999).