ITPedia is a public website focused on documenting projects by the people that make them without explicit permission by the community it represents. ITPedia is being designed using wiki software as it's foundation, which provides an informal and unstructured collaborative environment. The implementation will include a basic data structure using a loose semantic ontology as a way to overcome the inherent taxonomical weaknesses of Wiki software.

There is one design challenge in knowledge management (KM) efforts by which all other design challenges pale in comparison: "How do you get people to use the system?" KM became an established discipline in the mid 90s, relying on insight from organizational psychology, business administration, and information sciences. Most traditional KM systems are centered around a linear workflow approach, a system of document creation and approval into strict taxonomies that can be retrieved and understood by other users that had nothing to do with the original document. This highly structured approach usually provides very clear entry points for users, with tight mappings between pre existing responsibilities and defined repositories of knowledge. There are problems with this kind of approach, the primary being that people don't work in a linear fashion; we work in somewhat chaotic patterns that rarely resemble a linear workflow. Examining the changes tracked by a Microsoft Word Document is proof enough of this. Unless you can see how a group of users got to the final representation of their work, how useful can it actually be? The final result of a traditional KM system is that it languishes, many times viscerally hated by it's intended users while they go about making and creating elsewhere, great opportunity lost.

On the other side of the spectrum, Wiki software is a completely free-form system that has proven successful in organizing communities around collaborative efforts. Although the premise of it's strength is a lack of formal

structure, in the absence of user continuity it becomes a weakness. Wikipedia is obviously the most successful example of a community wiki, but it's defining characteristic is a purely voluntary and virtual nature. Although there are many lessons to be learned in terms of motivation from Wikipedia, it can also be deceiving as a model for a physical community. The specific community for ITPedia is ITP. ITP consists primarily of two groups of people at any given moment: people active on the floor and people inactive, previously active on the floor. The former are the ones that define the culture and the creation of work at any given moment, the latter are the ones who did exactly that in the past, but are no longer involved. Superficially, these two groups are students and alumni, but once again providing proof of taxonomical futility, it's common for an alumnus to come back and teach at ITP or a staff member to enroll as a student. the Active group is bound by a physical reality, the 4th floor of the Tisch building, while the inactive group resembles more a virtual community with a common physical past. Any set of informal rules and traditions created by a given active group at ITP are subject to be forgotten, ignored, re-written, and re-invented. The reason is that the core of the creative force, the students, have a 50% turnover rate every year. Although it's the nature of any educational system that some of the wheel must be reinvented in the process of learning, almost everyone would agree less, rather than more, re-invention is better. Thus, the design challenge rephrased: "How do you design a knowledge management system that structurally reflects how people work without becoming overbearing and rigid?"

The solution for this problem is a "slightly semantic" approach. The basis of the implementation will consist of a Mediawiki installation with semantic and form extensions. The first extension allows definitions of relationships between data in the form of a triple. A triple is a statement about a resource in the form of *subject-predicate-object*. A simple example is The sky (subject) has the color (predicate) blue (object). I'll spare the painful details of the semantic web dreams and just

say that a triple can be very useful if you don't pin all the hopes of human computer interaction on it. The specific reasons it is useful in our case is that triple relationships allow for the emergence of patterns and super categories, as long as the right kind of relationships have been defined and made accessible to users.

The right of kind of relationships are the ones that already exist between people, their work, and the resources they use. ITP is project oriented since so much of what we do (and don't do) is centered around our projects. The timeline, the url, the people who worked on it, the skill involved, the show it was in, the threads and links between different parts of the schema are the projects as much as the actual project itself. The more tangible of those links are the the people who make the projects: people cast longer semantic shadows than projects. The two primary views on ITPedia will be individual people pages and project pages. Most other pages will act as glue between those two. My central hypothesis is that most ITP projects are iterative across people and time. Our ability to track our intellectual roots within our own community makes our work mindful and better informed. Our ability to submit that work into a system respectful of the complexity of iteration and allow it to become part of someone else's intellectual roots makes us humble creators. Participating in something something bigger gives us a strong sense of satisfaction and pride.

My other hypothesis is that provided a basic semantically minded structure, the highly enterprising user-base will find interesting ways to use those predefined relationships to create other kinds of structures; this is where having a wiki as the underlying foundation is a huge advantage. For example, in the current schema of predefined templates (people, project, classes) there is no template for a list of restaurants around NYU. It would take the addition of one extra property, "cuisine" and the creation of a simple form to create that kind of addition to the

schema. This is exactly the kind of flexibility that allows a knowledge system to adapt organically to the needs of it's user-base, instead of slowly rusting away. By establishing a set of properties that provide an organizational equivalent of the primary levels of Maslow's Hierarchy of needs, users themselves can become familiar over time with the most basic of metaphors and begin to extend them in the way that fits their own higher level patterns of behavior.

The target audience as previously mentioned is the ITP community at large. Within that community there are multiple groups that overlap, with different agendas. Below are the groups identified and their desired behaviors.

1st year students. The first step is to establish a familiarity with ITPedia and what it can do. In simple terms, first year students should be able to transform their behaviors from "How do you use a wiki" to "that information is on the wiki" within the first two weeks, and "That information should be on the wiki, I'll put it on there" within the first two months.

2nd years students. In terms of continuity, the transfer of tradition and best patterns is the most important behavioral aspect. They should be able to identify their own cultural replacements among the 1st years and mentor them into taking over similar ITPedia roles. 2nd years should be motivated by a sense of legacy to make sure all of their projects have proper representation in the knowledge base.

Full time/adjunct faculty. The primary desired behavior of the faculty is to maintain a full syllabus for each of their classes on ITPedia. In addition, the full time faculty would ideally tend to cumulative area of focus pages that fall under their expertise. For example, the faculty member working most in the realm of mobile computing would tend to the general mobile computing area page in a curatorial fashion.

TNO/Drive-by Kings. The students in charge of these two traditions already use the current student wiki as an organizational tool. With ITPedia, they'll have the ability to augment those lists with additional information. For example, each TNO outing would ideally have links to outside photosites with pictures from that night. Since those links show up on the list, it would be trivial to add those to their corresponding page.

The technical implementation of ITPedia is divided into two sections. The first is the infrastructure and seeding which is on it's way. Instead of presenting a blank slate to the community, ITPedia will import XML feeds of all the existing datasets in the community that are accessible. Fortunately, work is being done in creating RSS formatted / machine readable outputs of the major databases at ITP. Specifically, the Student Roster, The projects, venues, classes and ER equipment databases are all available. As a result ITPedia will kick off with a well seeded dataset. This will be the first time all of these databases will be brought together in one single flat scheme in a way that is accessible and editable by any of the students, alumni, faculty and staff in an equal manner. The high level of transparency and exposure should provide users with a strong motivating factor. Some examples would be editing personal profiles to include up-to-date information, linking person pages with appropriate work and editing descriptions of projects for quality. The basic empty structure should be done by march 7th, and the data from the databases parsed and populated using the pywikipediabot scripts by March 21st.

The second aspect of the implementation is how to increase participation through interventions and awareness. The first proposed strategy is focused around creating hooks into other tools of communication and collaboration currently used in the community. The first that comes to mind is the student list, (which although

popular - 58 active threads in the last 24 hours - is not used by everybody at ITP). Automating emails from ITPedia to the student list on a weekly basis with items such as 'New Pages', 'Updated Pages', 'Most Active Users', 'New Categories' would serve to provide an constant reminder of how ITPedia is being used and by whom. Within that lies the potential for less then regular contributors to see the potential for how they can use ITPedia in ways they haven't previously. in addition, the 'Active User' list can invoke a sense of pride in a community that attaches a high level of recognition to members that tend to the community itself. The other tool in this category would be creating better ways of invoking emails for pages that people are interested in. This kind of email notification would be both for community pages users have created and have an interest in specifically, but also group collaboration pages that notify when a different member of the group makes changes. This is a very important point. If ITPedia is ever to become the kind of knowledge management system that transcends the linear pitfalls, it needs to give it's users the ability to use it from the first moment of a brainstorm session. This way a finished project from the ITP show could be traced all the way back to it's first unorganized document created in the second week of class. Currently this kind of work ends up living in google documents shared by 2-3 people. The nature of this kind of work is of low importance to the maker in terms of final documentation, and as a result would never get posted to an ITPedia page if it wasn't already created there. The solution to this is a suite of better notification tools, 'link beachheads' in the class ITPedia page template for the creation of group projects, and more useable WYIWYG interfaces for the collaboration pages. Since the competition is Google, the bar is set pretty high. However, considering the enticing hooks of syllabus and class pages, not impossible. The vast majority of this work can be done with mediawiki extensions and fine-tuning of templates, scheduled for the end of March.

The next implementation scheme involves creating physical awareness of ITPedia on the floor. Since this is an online representation of a physical community, it will live or die by how aware the current student base is of the site, no matter how many alumni or faculty contribute to it. The first step is to simply set up a computer in a central location with a huge sign that says ITPedia. Not high tech, but considering how successful projects like "Social Genius" have been in generating community interest, it should at least make some headway into the portion of the student body that doesn't seem interested in list participation. The second step is a multimedia version of blogblender that displays a boxee setup on the floor of everyones flickr/vimeo/youtube accounts. A setup such as this one would pull in the rss feeds in ITPedia people pages and set them up for consumption. Again, blog-blender has proven that a beachhead of accumulated community syndications does a good job of cross breeding ideas and references.

Between the two proposed technical approaches, I feel the first one is more important and relevant to the primary needs of the system. As much as creating a higher level of on-floor awareness is important by displaying the information gathered, if the tools don't exist to allow a full range of qualitative entries from the largest number of people, then there simply won't be enough to show on the floor awareness strategy. Although the basis of ITPedia is a machine readable structure from different databases, this cannot be enough to ensure success. Success for a project like this is true character emerging, and that simply doesn't reside in 3 sentence descriptions and semantic triples. That stuff is just the groundwork. That character resides in anecdotes, ad-hoc group activities, and gossip that must find it's way into a permanent home that future students and the outside world can see. If a first year said on her first day "Sure, I know why the japanese room is called that, I read it on ITPedia when I was getting my

application ready", as trivial as that may seem, that will the proof of a successful knowledge management project.