The Digital StudyHall

The DSH Team
dsh@cs.washington.edu

The shortage of qualified teachers in rural and slum schools is one of the most challenging problems faced by over-burdened education systems such as the one in India. It is not unusual for a government village school housing 200-500 students to have fewer than half a dozen teachers, who are not well-qualified to teach many of the subjects that they are required to teach. And it is not unusual for quite a few of even these teachers to be absent due to various reasons on any given day. In the Digital StudyHall project, we are building a system that seeks to help kids and teachers from such schools.

To avoid retracing the missteps of earlier "wire-the-schools" or "distribute-computers" projects, we follow two important principles: (1) cost realism, essential if we are to scale the system up to a significant number of schools and students; and (2) building systems that solve end-to-end education problems, so the twin pillars of technology and pedagogy must develop side by side, and content, people and organizational relationships must play leading roles.

We discuss the following aspects of the Digital StudyHall system (DSH): (1) a comprehensive digital video database of K-12 curricula in local languages; (2) an architecture of a network of hubs and spokes for scalable and locally relevant content production, teacher training, monitoring, and evaluation; (3) mediation-based pedagogy that puts local teachers and student peers in the drivers' seats and allows the teachers to be trained in an ongoing and specific manner; (4) A practical infrastructure for sharing community-generated video; (5) A "hybrid" network integrating SMS, voice, localized radio broadcasting, the postal system, TVs, video-players with other connectivity technologies for reaching the vast under-served population. We close the write-up with a brief status report and some conclusions.

1. A "People's Database of Everything"

The Google book scanning project is attempting to create a digital database of all books ever written by humans. In a similar spirit, in the Digital StudyHall project, we are working on creating a digital *video* database of "everything." (Figure 1 shows one page from the DSH database.) The DSH database, however, has some important unique characteristics, compared to the book-scanning project or any other existing educational content creation projects.

First, content creation in DSH is a community-based effort, as opposed to the traditional model of a small number of "experts" holed up in an office far removed from the target audience. (Think the difference between YouTube and Fox News.) This approach has important scalability and local relevance implications. The grassroots contributors to the DSH database include best teachers in middle-class urban schools, best teachers in rural schools, students, and other idealistic and enthusiastic volunteers such as retired university professors, scientists working in government labs, college students overseas, and various NGO staff members. In short, the database is created by the people, and for the people.

Like Wikipedia, this people's database should become an open-source and continuously evolving community resource. Unlike the text-based and largely unstructured nature of today's Wikipedia, however, the DSH database is video-centric and highly structured. This second crucial characteristic of the DSH database, its video-centric nature, is critically important. India, like most developing nations, is a predominantly rural country, and suffers from a low literacy rate. (By optimistic estimates, the adult literacy rate of India is about 60%.) The immediacy and

liveliness of videos, shot in local dialects, of teachers, experts or peers speaking about and acting on issues that the audience can relate to and care about, is likely to have a far more profound impact than, for example, a book-scanning project predominantly focused on texts written by dead white males in a foreign language.

Figure 1: A page from the DSH database.

This video-centric approach has other important advantages compared to traditional forms of educational content creation, which is typified by authoring flash-ware and slide-ware. A quick glance at any accomplished teacher should convince even the most casual observer that it is difficult (if not impossible) to encapsulate what a good teacher can accomplish in a piece of flash-ware. Good teachers are often the best "performers." It is their "showmanship," their way of relating to their audience, their people skills, that are their most potent tool, a tool that is beyond the reach of flash-ware or slide-ware. Video, though not perfect, comes closest to capturing the "performance" in flesh. Another problem with the flash-ware- and slide-ware-only approach is that it is slow and expensive. One must insert a technical professional into the authoring process, who not only introduces an additional element of disconnect between the producer (the teacher) and the consumer (the students), but is also a scalability bottleneck (in terms of cost and skill at least): it is simply impossible to build a database of "everything" quickly and cheaply based on a flash-ware- or slide-ware-only approach. Video is one of the few means that enables a large fraction of the population to become effective authors.

The third important characteristic of the DSH database is that, at least at the beginning, its content is predominantly based on systematic and coherent sequences of lessons derived from curricula designed and sanctioned by state government boards of India. The video recordings adhere to the curriculum topic by topic, lesson by lesson, and chapter by chapter. This is in contrast to unstructured and nearly random collections of snippets of knowledge found elsewhere that is difficult to find easy acceptance or adoption in existing schools, which by law or by choice, overwhelmingly tend to strictly follow official syllabus. The fact that our video

recordings follow official boards, however, does not conflict with the need for innovation. Indeed, the model teachers featured in the database tend to embellish well beyond what is contained in the textbooks; their lessons are typically highly interactive and activity-based; they incorporate established educational principles in these daily lessons; and we produce a stream of complementary materials, such as digital stories, science courseware, drama, student projects, and other materials to further complement the content dictated by the syllabus. (Some of these complementary materials, however, are produced using slide-ware and flash-ware and then subsequently turned into videos---the limitations of the slide-ware- and flash-ware-based approach do not pose a handicap in this case because its role is complementary.)

The above three key characteristics, the community-based approach, the video-centric nature, and adherence to official curricula (initially), directly lead to our goal of building a "people's database of everything." The official curricula provide an initial broad definition of what we mean by "everything," and the community-based and video-centric approach provides the only feasible path: without the scalability afforded by it, building a database of "everything" is merely an empty slogan. We envision building a database that contains *all* subjects, encompassing *all* grades, spoken in *all* local languages, covering *all* state and national boards. (We are also beginning to apply the same system to agricultural extension and frontline health worker training.) This is a database that every child (and adult, for that matter) should have access to. We believe such a database may have profound implications for liberating knowledge, democratizing learning, and revolutionizing education.

2. A Network of Hubs and Spokes

DSH is not a physically centralized system. Instead, DSH is designed to work as a decentralized network of hubs and spokes. Each hub is a center of education excellence and the hubs themselves communicate with each other. The spokes are typically the under-served rural and urban slum schools that need help the most, schools that lack good teachers, good content, and other resources. Each hub works on content production (typically in a local language), content dissemination in its neighborhood, teacher training, monitoring, and evaluation, and interacting and sharing with other hubs. The hubs-and-spokes model is how we may effectively scale up the DSH system.

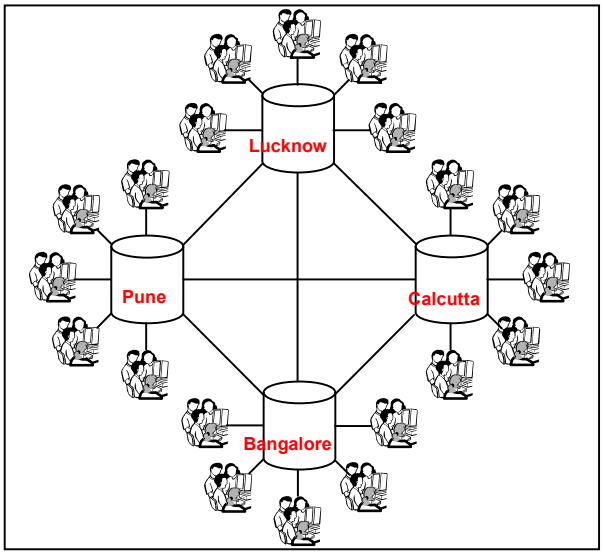

Figure 2: The distributed hubs-and-spokes model.

In addition to being a scaling vehicle, another important role served by the hubs-and-spokes model is what we call "impedance matching." This term refers to the need of ensuring that the content being produced at each hub is appropriate for the audience in its neighborhood. The hubs are frequently co-located with good urban schools that target a middle-class student population. (This is the case for both of our Lucknow and Pune hubs.) In one of our earliest attempts at the Lucknow hub, we made recordings of regular classes in such a school and learned that such content was not appropriate for the target audience in under-served schools. Factors such as language (English-medium instruction in rich schools vs. Hindi-medium elsewhere), syllabus (CBSE national board vs. UP state board), and student background differences contributed to a big gap between middle class urban schools and under-served schools, and made a direct transfer of content between them ineffective.

After several experiments, we settled on a "hybrid model" that was intended to bridge this gap (Figure 3). We recruited the best teachers from the middle class schools, but instead of filming their regular classes in front of the middle-class students, we staged specially designed classes for poor girls from the neighboring slums. The best teachers know to adapt to their audience and adjust their style and content accordingly, so this hybrid model combines the best of both worlds: top-quality teachers and an appropriate student audience, allowing us to produce content and experiment with pedagogy that is meaningful for our target audience in the spoke schools in rural and slum areas. For a similar purpose, at our Pune hub, selected teachers from a middle-class school are sent to teach and get recorded in poor government schools in the neighborhood.

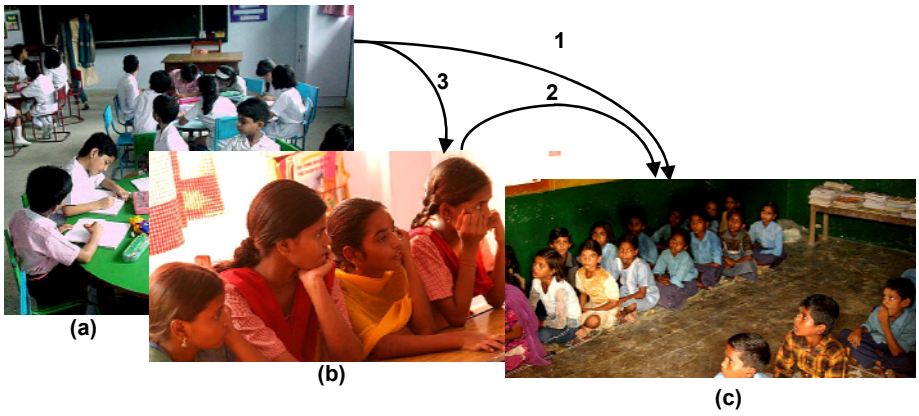

Figure 3: A tale of three schools. (a) The urban StudyHall (middle-class) school. (b) The afternoon school for girls from slums, the Prerna school, housed in the same premises of the middle-class school. (c) The Mandantoosi village school. The direct transfer of content from school (a) to school (c), labeled as arrow (1), does not work well. Instead, we use a combination of arrow (3) and arrow (2), as we film the middle-class teachers in front of the slum school audience, and transmit the resulting content to other poor schools.

Yet another way for a hub to accomplish scalable content production and ensuring its relevance is to involve the under-served spoke schools themselves in the content production process. Under this approach, explored in our Bangalore hub today, we identify the best teachers in the spoke schools served by the hub, organize them in a regular recording schedule, and the resulting content is shared with the other peer spoke schools. This approach not only ensures content relevance, but also motivates the spoke school teachers and allows them to gain confidence: the teachers being recorded strive to learn and use the best methodology to put on the best shows they can, and the peer teachers who receive the content are inspired to match their filmed peers. Involving the spoke school teachers in this manner is perhaps an even truer manifestation of the philosophy behind the "People's Database."

3. Mediation-Based Pedagogy

The principal means of disseminating the content in the DSH database is shipping DVDs to spoke schools. Each spoke school is given at least a TV and a DVD player. (Most schools also need a big lead-acid battery and DC-to-AC inverter for dealing with intermittent electricity. We are also working with engineers to experiment with a bike pedal-powered generator and other approaches for locations that lack access to the electricity grid altogether.) An obvious question is whether kids can learn by just watching TV. The short answer is no. The long answer starts with the note that we need local teachers and accomplished peers to play an active role in "mediating" the content shown on TV.

Put simply, "mediation-based pedagogy" refers to the need of placing a teacher (or a "mediator") in between the students and the TV. The mediator periodically pauses the video and engages the students in various activities based on what has just occurred on TV. These activities may include asking questions, inviting kids to do board work, and organizing role-playing activities. The mediator's job is to make his or her class as lively, dynamic, and interactive as the one conducted by the model teacher on TV. In effect, the video and the mediator form a "team:" the video provides an example, a framework, a lesson plan, and a content and methodology model; while the mediator, who may not be highly skilled in some domain-specific knowledge, supplies the crucial interactive element. Usually, neither half is sufficient by itself, but an effective combination of the two can become a powerful instrument for addressing the critical problem we face: the acute shortage of trained and qualified teachers, especially in rural and slum schools. (Features such as allowing the mediator to play and pause the video in a way he or she sees fit, at a time and place of his/her own choosing, and the need of delivering highly customized local content, are what makes the DSH system more similar to an on-demand Internet-based video sharing system than a rigidly scheduled and one-size-fits-all traditional broadcast TV that does not allow any of these flexibilities.)

While placing the local village school or slum school teacher as a mediator in between a TV and the students (Figure 4(a)) may be the most obvious choice, there are several other ways the mediation can work. One is "peer-mediation," the approach of recruiting the brightest fellow students to serve as mediators during periods when the local teachers are absent (Figure 4(b)), which are common occurrences in government schools in India. In our experience, the student mediators appear to universally display a high degree of responsibility and enthusiasm when they are put in charge, and a different social dynamic of peer-mediation can play an effective complementary role to the other instruction received by the students.

(a) (b)

Figure 4: Mediation-based pedagogy. (a) A village teacher mediates a live class. (b) A peer-student mediates for her classmates.

Yet a third way for the content to have a positive impact is for the local teachers to study the supplied videos on their own, ahead of the live classes. This typically happens with teachers who are highly motivated and want to improve themselves. Without DSH, such motivated teachers do not have an effective means of self-improvement; with DSH, a path is open to them as long as they are willing to work hard. This path leads to a best-case scenario, when we end up "graduating" teachers, who may indeed choose to cast aside the live mediation crutch and become effective teachers in their own right in the long run.

Regardless whether the local teachers use the DSH content to train themselves ahead of classes or to mediate live classes, they end up training themselves to be more effective teachers. There are three levels of skills involved. The first set of skills concern being an effective mediator. The individual DSH hubs are responsible for running live face-to-face training workshops, attended by the spoke teachers (as well as some of the model teachers appearing in the videos). Such a training workshop can be completed in a couple of days. The second set of skills concern broader pedagogy. Parts of these skills are also covered by the face-to-face training workshops. The videos that the spoke school teachers carry home for their daily use, however, serve as a continuation of the training workshops, so the spoke school teachers get to observe and study on a daily basis how the best teachers teach. The third set of skills concern domain-specific knowledge, such as that is required to be a fluent English speaker for an English teacher or math knowledge for a math teacher. The spoke school teachers get to improve their own knowledge on these subject matters through their daily use of the videos. In traditional teacher training workshops that last just a few days, the short duration necessitates that the topics covered must be kept at an abstract level, and it is not always clear how such abstract principles should relate to the daily topics. In DSH, the videos carried home by the participating teachers provide an ongoing and highly specific training, so this mode of training has the potential of being much more effective.

In short, the focus of DSH is not to replace people with technology; instead, it is about amplifying the reach and the power of the relatively small number of the skilled teachers, and to train and empower the less skilled teachers. In this sense, DSH is foremost a "people system," not just a computer- or network-system.

4. Technology for Sharing Community-Generated Video

The e-learning landscape is littered with misguided and expensive "wire-the-school" projects that have little to show for in the end. To avoid retracing those missteps, one of the most important principles that we follow is cost realism. A conventional "wire-the-school" attempt is simply not feasible for large-scale replication in rural India today or in the near future. Technically, what we have described so far is akin to a user-generated video sharing system. The question we face is how to build such a "Web 2.0" application without having to replicate the "Web 2.0" physical infrastructure in slums and rural areas today.

A best example illustrating our approach is what we call the *Postmanet*, in which computer network packets normally placed on wires are now placed on DVDs transported by the postal system. The Postmanet allows us to have pervasive, high-bandwidth, and low-cost asynchronous connectivity to just about any place, including the remotest areas. (Such a high-latency and high-bandwidth channel can also be complemented by a low-latency and low-bandwidth link, such as a cellular or a packet radio link.) On top of this low-level connectivity provided, we build the rest of the distributed DSH database in a way that is conceptually similar to how an existing peer-to-peer content sharing network works: there exists a single coherent DSH database name space that all the distributed sites can read from or write to (Figure 5).

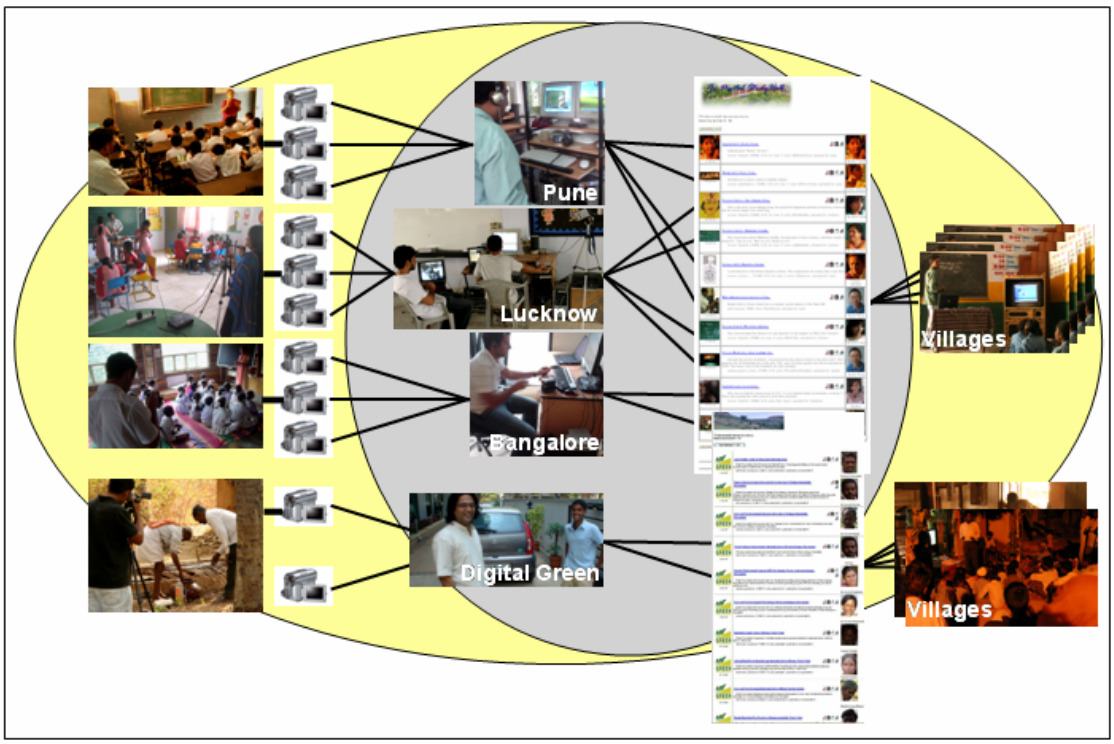

Figure 5: A peer-to-peer system for sharing community-generated video.

While the description above provides unconventional interconnect and database solutions, we also need to solve "input" and "output" problems for the system. As we have discussed earlier in Section 1, keyboard-only or text-only input is not sufficient for a literacy-challenged society; and we cannot afford a large-scale acquisition and deployment of PCs to serve as input- and output-drivers. Figure 5 illustrates our end-to-end system. On the far left is a fleet of digital camcorders serving as the eyes and ears (or input devices) of the system. They are shared among the DSH participants who contribute content into the system and, like rental cars, they may be constantly on the go. The resulting tapes are funneled to the nearby hubs for digitizing and uploading into the local databases, which communicate with other instances of the local databases at other hubs via the Postmanet and synchronize their content. On the far right, shared TVs and DVD players economically serve as the output devices of the system.

It is useful to compare the Digital StudyHall infrastructure against satellite-based approaches. Satellite-based approaches are expensive and they require a great deal of support infrastructure. Satellites are a good broadcast medium: a small number of one-way streams consumed by a vast number of content consumers. But broadcast models are poor ways of delivering customized content and allowing two-way exchanges. Satellites can also be used to support non-broadcast or even two-way communication. If we do that, however, we face a severe bandwidth problem: each of a large number of communication channels only gets a small fraction of the aggregate bandwidth. The bandwidth limitation is especially serious on the uplinks. One important advantage of the Digital StudyHall approach is that it allows high-bandwidth, any-to-any, point-to-point communication, which in turn enables a high degree of content customization and rich two-way exchanges.

In addition to being a technical mismatch, satellites are also a poor pedagogy match. Satellite-based communication was originally intended for synchronous real-time applications such as live teleconferencing. What we have observed is that teleconferencing is a poor metaphor

for enabling distance learning in rural areas: when an experienced remote teacher attempts to "teach" to a large number of rural sites over a live teleconferencing-style link, the many students at the many target schools cannot "pause" the remote teacher or participate in any meaningful interaction across distance on a large scale. The synchronous nature of the communication, paid for at a hefty price, ends up being a liability. The asynchronous nature of the DSH communication, in terms of DVDs that allow local teachers to take full control, to play and pause the content at will, at times and paces of their own choosing, to engage in meaningful dialogs with their local students, to train themselves after school hours, turns out to be a blessing.

5. A Hybrid Network for Reaching a Vast Under-Served Population

In the conventional "Web 2.0" that we have grown to love, all those who participate in the system as content contributors and consumers must acquire computers and broadband connectivity in order to "play." Instead of relying on this one-size-fits-all model, we at DSH are exploring a more practical "two-tier" model. Again, consider the illustration in Figure 5. In the center is a "high-tech core," namely, the hardware, software, and human expertise elements employed inside the DSH hubs. These include high-end PCs for video processing and database management, significant amounts of storage, video digitizers, DVD-burning robots, broadband connectivity, elements of a distributed database, a wide variety of video and audio processing software, and a relatively technology-savvy staff managing the core and working closely with domain-specific experts in education, agriculture, and healthcare. The "core" is akin to "the cloud," in Internet jargon; one difference is that the core here is embedded inside and run by the community.

Outside this "high-tech core," there lies a "light-tech fringe," which employs more practical, simpler, cheaper, and better-understood metaphors, including TVs, DVD players, camcorders, phones, and other new devices that we are working on. While these devices may begin with humble origins today, over time, more intelligence and sophistication may creep in to deliver better functions. For example, in the near future, one may be able to plug a TV, a DVD player, and a cell phone together; browse a media-rich catalogue contained on a DVD delivered by the postal system; issue orders via SMS over the cell phone link; and receive more discs in the future. Regardless the level of sophistication employed in this "light-tech fringe," the point of this two-tier arrangement is to make a "Web 2.0"-like system accessible to the vast underprivileged population, so that a much bigger segment of the society can contribute content to the "network" and consume content made available on it.

This hybrid two-tier network approach is not only applicable to community-based video production and sharing; it can also be used for other application scenarios. For example, we are currently building a hybrid network integrating community radios, voice over land-line and cell phones, and SMS messaging. An example application is a voice chat system that allows students and teachers from multiple DSH villages to chat across a long distance: participants can send "input" by simply placing cell phone calls to a hub server, and receive "output" as they hear the entire community chatting on their regular FM radio receivers, which receive their signal from a local FM transmitter, which is in turn driven by a local village "base station," which in turn communicates with the hub via cell phone voice and SMS traffic. Teachers, staff from the hub, and even volunteers from the US may participate in such a "chat room," conducting a wide variety of educational activities, such as science questions and answers, story and book reading, math quiz and game shows, complementing in-class curricula in such after-school activities. These chat sessions can also be digitally stored, spliced, and re-used in the future. In this scenario, again, the "fringe" is populated by simple, cheap, and practical devices like conventional FM radio receivers and cell phones, while the "core" contains the intelligence, storage, and communication means to drive these fringe devices, uniting them into a coherent

whole, providing functionality above and beyond what these fringe devices are capable of in isolation today.

What is interesting in this two-tier system is *not* a *hardware* story based on this sudden revelation that somehow cheap and simple devices at the fringe are good and can automatically save the world; what is key is an ongoing *software* story: it is the software that will knit these devices into a coherent fabric that behaves more like a network. For example, it is the software that makes an old beaten up FM radio, or a conventional phone, or even SMS messaging participate in a networked "chat;" it is the software that makes these devices come alive and play new roles that they were probably never intended for. This emerging software infrastructure is only beginning to take shape. And this developing software system is in no way necessarily tied to the lowest common denominator of low-end devices...

In this section, we have discussed ideas of employing a practical hybrid network to reach the under-served population. The higher-level DSH goals, such as producing a People's Database of Everything freely accessible to all, and harvesting community participation, however, are technology-neutral and target audience-neutral. So, for example, in countries where there is a sizeable middle-class population today, the DSH database may be accessible via PCs, broadband, and/or portable devices (like the Zune or the iPod); a DSH mediator may even play her role remotely over the network. When, in the future, technology options improve in poorer regions, we may also provide a smooth migration path, transitioning gradually from "light-tech" delivery to "higher-tech" delivery, and improving service quality in the process. The same system may be used for other areas of work (such as agriculture, health care, and political activism). We hope DSH will play a role in transforming the education landscape, spanning a continuum across the dimensions of space (from the disadvantaged to the middle-class) and time (from now into the future).

6. Status and Conclusion

A live deployment of the DSH prototype has been in experimental use by students and teachers in and around the Lucknow hub starting in July of 2005. Starting in the summer of 2006, we have launched two more experimental hubs in Bangalore and in Pune. In the period of about two years, we have accumulated about 600 high-quality recordings of lessons staged by the best teachers at the hubs, and about 400 other pieces of content, which include science courseware, digital stories, recordings of drama performances, and other student projects. The languages used in the content include Hindi, Kannada, Tamil, Marathi, and English. As the content is quickly and cheaply generated, it is being continuously pushed out to about thirty rural and urban slum pilot schools around the hubs. Preliminary results appear promising; and we are currently in the process of performing more rigorous and systematic evaluations.

In this paper, we have described the Digital StudyHall system, a system that strives to deliver the best instruction digitally to rural and urban slum children in a highly cost-effective manner. The aim is never about replacing people with technology; instead, DSH is about facilitating communication. A key to the success of this project is to proactively engage both students and local teachers in a continuous dialogue so they learn by communicating, participating and doing. We hope to eventually scale up the system to cover a far greater number of villages and children, contributing toward the Millennium Development Goal of universal primary education.

For more information about the project, please visit http://dsh.cs.washington.edu or email dsh@cs.washington.edu.