

The Digital StudyHall

Randolph Wang* Urvashi Sahni[†] Sumeet Sobti* Nitin Garg*
Jaswinder Pal Singh* Matthew Kam[‡]
Arvind Krishnamurthy[§] Thomas Anderson[¶]

Abstract

In this paper, we describe a distance learning system that would allow resource-starved village schools in rural India to benefit from the better human and content resources available in the urban environments. The e-learning landscape is littered with misguided and expensive “wire-the-schools” projects that have little to show for in the end. To avoid retracing these missteps, we must follow at least two important principles in our solution: (1) cost realism, which is essential if we were to scale up the system to encompass a large number of villages, schools, and students in the long run; and (2) building systems that solve end-to-end education problems, instead of narrowly focusing on just providing connectivity.

The proposed Digital StudyHall system has the following novel key components. The first is a generic digital communication mechanism that places bits on storage media transported by the postal system instead of wires. This mechanism, the *Postmanet*, provides pervasive, high-bandwidth, and low-cost asynchronous connectivity to just about any place. When combined with a low-latency channel, such as a packet radio connection, we may combine the latency and bandwidth advantages of both channels. Robotic arm-based automation in our headquarters site further enhances transparency and efficiency. The second is a mechanism that turns regular TV screens into “networked thin client displays.”

This mechanism, which we call *EdTV*, lowers the cost of end user devices, and truly bridges the last mile by leveraging TV and radio control signals. The third is a web repository that collects education content, and connects learners and teaching staff across time and space, so staff in urban schools and volunteers (potentially from overseas) can contribute in a way that allows them to make flexible time and location commitments. This site, dubbed the *learning eBay*, would be accessible via both conventional networks and the Postmanet. These components would enable a wide variety of digital education “workflows,” such as lecture capture and replay, homework collection and feedback, and question-answer sessions. We also plan to perform pedagogy research on the Digital StudyHall, so that it can serve as an effective learning science testbed, tightly combining education research and practice.

1 Executive Summary

In the past decade, Dr. Urvashi Sahni, a native of Lucknow, India, and a co-author of this paper, has been working on improving basic education for disadvantaged children in rural areas and urban slums in her native state. “StudyHall,” a highly regarded school run by Dr. Sahni in the city of Lucknow, caters to both middle-income students (in regular classes) and girls from the urban slums (in an after-school program). StudyHall also subsidizes six affiliated village schools. Like the vast majority of typical rural Indian schools, these affiliated village schools experience a severe shortage of well-trained teachers.

We would like to build on Dr. Sahni’s previous work, and construct a “Digital StudyHall” that allows resource-starved village schools to benefit from the better human and content resources available in the urban environments. As an initial step, we would like to knit the StudyHall headquarters and the affil-

*Department of Computer Science, Princeton University, {rywang,sobti,nitin,jps}@cs.princeton.edu.

[†]The StudyHall Foundation, urvashisahni@yahoo.co.uk.

[‡]Department of Electrical Engineering and Computer Science, University of California Berkeley, matkam@cs.berkeley.edu.

[§]Department of Computer Science, Yale University, arvind@cs.yale.edu.

[¶]Department of Computer Science and Engineering, University of Washington, tom@cs.washington.edu.

iated village schools into a networked cohesive whole so they can better communicate and systematically share content with each other, so, for example, high-quality lectures digitally captured at the headquarters are made available to the village schools, and homework and questions from the village schools can be handled by the staff or the volunteers at the headquarters.

In the longer run, our goal is to scale up the Digital StudyHall to serve a much greater number of rural students, including those living at places that have no school today, and to allow qualified volunteer teachers beyond those in Lucknow, including volunteers overseas, to contribute remotely via the system. Note that our goal is not to compete against or replace human teachers; on the contrary, our goal is to amplify the reach and power of the limited number of qualified teachers and volunteers that we do have.

1.1 Guiding Principles

The e-learning landscape is littered with misguided and expensive “wire-the-schools” projects that have little to show for in the end. We do not intend to retrace those missteps. There are a couple of principles that we intend our project to follow. The first is *cost realism*. Consider a recent survey of schools in Bihar, Madhya Pradesh, Uttar Pradesh, and Rajasthan [10], which shows that 63% of the schools have leaking roofs, 58% have no drinking water, 89% have no functioning toilet, 27% have no blackboards, and 8% have none of the above! The high cost of a large-scale conventional “wire-the-schools” attempt must be carefully weighed against these pressing basic needs. Priority-setting and cost/benefit analysis is crucial [14]. Keeping the per-school expenditure down would allow us to eventually scale up our system to encompass a greater number of schools and students.

The second principle is focus on problem solving and *system building*. Providing connectivity is only a part of a solution. Consider the much-touted launch of EDUSAT, India’s first educational satellite placed in orbit in September of 2004. Five months after the launch, all the pieces are functioning in space, but there is virtually no ground support to enable the delivery of distance education to remote parts of rural India, which was to be the satellite’s chief benefit. Nor are the technology, the content or the trained manpower needed to allow the satellite to explore its full potential [8]. After a massive initial investment,

the maximum that can be achieved this year is adding to existing learning facilities at a handful of established national universities. Given the mountainous support tasks that still need to be accomplished, and the fact that EDUSAT’s lifespan is only seven years, it appears that India is unlikely to be able to take full advantage of the money and effort it has invested in developing and launching the satellite. The lesson is that we need to focus on building whole “systems” that solve end-to-end problems.

1.2 Innovative Claims

The proposed Digital StudyHall is built on top of the following three novel key components.

- *The Postmanet* exploits digital storage media (such as DVDs) transported by the postal system as a generic digital communication mechanism. In essence, under this approach, network packets normally placed on wires are now placed on DVDs instead. While the idea of sending digital content via the postal system is not a new one, none of the existing attempts (such as Netflix) have turned the postal system into a truly *generic* and *transparent* communication channel—to fully realize the potential of this approach, and to transparently scale up the system, one essentially needs the analogous equivalent of a networking software stack (such as TCP/IP) that manages point-to-point as well as end-to-end transmission issues (such as the analogy of a “packet loss”).

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 link or a packet radio link, where it is available, so we can combine their bandwidth and latency advantages.

In the Digital StudyHall, a key application of the Postmanet is a Postmanet-based http (or *phhttp*) mechanism that allows village schools to interact with a web-based content repository housed at the Lucknow headquarters. Under this approach, http requests, replies, and server script fragments can be entirely encapsulated on DVDs transported by the postal system. A robotic arm-operated DVD processor at the headquarters allows us to automate almost all aspects of the phhttp operations.

- *EdTV* allows us to turn regular TVs into “thin

client displays” driven by a shared computer. The most immediate benefit of this approach is that it provides a low-cost solution to the “display problem” in village school classrooms: the problem of not having enough computer displays for all students to see clearly.

The implication of the EdTV approach, however, reaches beyond the classroom, both literally and figuratively. Any TV sets within the transmission range of the inexpensive transmitter powered by the village computer can receive the computer’s display content, and a cheap custom-made radio transmitter allows a TV viewer to beam control commands back to the village computer.

This approach solves several difficult problems: it lowers the cost of end user devices, and it truly bridges the last mile by leveraging TV signals in one direction and radio control signals in the other. EdTV opens up a “portal” for interesting computer-based applications into a poor household that would have otherwise been beyond the reach of more conventional technologies. (We discuss some of the non-education applications near the end of this paper.)

In a way, EdTV is similar to kiosks. EdTV essentially brings a “face” of a kiosk onto each TV screen within the transmission range. This allows one to access services from the convenience of one’s home at a time of one’s choosing, and it allows multiple people in front of multiple TV screens to share a collective experience. These are the advantages that kiosks cannot provide.

- *The Learning eBay* is a site that connects learners and teaching staff across time and space; this is analogous to how an auction site matches supplies and demands. Volunteers and professionals all over the world can contribute content, conduct online teaching sessions, grade homework assignments, run virtual “office hours” via the site. Students and staff with limited experience or resource may tap into this site to augment their learning. Such a site would allow volunteers (potentially from overseas) to make flexible time and location commitments. This site (or the content repository) will be housed at the headquarters StudyHall school in Lucknow. People all over the world with conventional network access can access the site via the web. People who have no conventional network access, or those with limited bandwidth, including those in the village schools that we serve, access the site via the Postmanet-based http (or phttp) mechanism.

The repository will initially be rapidly populated with digital captures of high-quality lectures given at the headquarters school. Other types of content, which may require more time to develop, will be added gradually. All these materials will be made available to the village schools via phttp. Village schools lacking qualified graders may submit their homeworks to be graded and receive subsequent feedback digitally. We will also work with volunteers from a neighboring teacher training institute to improve the content and the services available through the site.

These components share a great deal of synergy. The combination of the repository and the Postmanet, namely the phttp-accessible repository, provides an abstraction that is akin to that of a distributed file system: it makes a single name space available “everywhere,” including those places that lack conventional networking access, and allows read, write, navigation, and other primitives to be performed on this name space. This file system analogy means that the phttp-accessible repository is a sufficiently general abstraction that can be used to support other applications. The combination of the Postmanet and EdTV also represents a natural two-hop solution: the ubiquitously available Postmanet feeds the EdTV “base-station” with high-bandwidth content at low cost, and EdTV completes the “final leg” by making available to the households cheap end viewing devices and two-way TV and control signals. And finally, the repository abstraction provides a convenient abstraction to build shared EdTV applications with. A common theme of these components is that they allow sharing and the delivery of highly customized content and experiences.

All these mechanisms are designed specifically to address the two previously mentioned guiding principles: low cost, and building end-to-end problem solving systems. For example, the Postmanet is a low-cost but highly effective network; EdTV is a low-cost but highly effective “networked thin client display;” and the homework feedback mechanism used by the learning eBay, which we discuss in greater detail later, minimizes printing so we do not have to spend too much on expensive printer cartridges and printer papers. We carefully “penny-pinch” on the amount of equipment each village school requires so we can scale up the system to cover more schools and kids. And instead of narrowly focusing on connectivity, the learning eBay is designed to be a problem-

solving *system* that addresses “workflows” such as those of content capture and delivery, homework feedback, and question-answer sessions.

Of course, the technologies are only half of the story of the Digital StudyHall—the other half is pedagogy research. Past experiences with comparable systems [6, 5, 1] indicate that instigating interactions with and among students can significantly improve the effectiveness of these digital teaching systems. The Digital StudyHall faces unprecedented challenges and opportunities. Many more types of interactions are possible: asynchronous digital interactions between villages and the headquarters, face-to-face interactions between the local village staff and the students, and those among students. A goal of the pedagogy research is to devise the most effective means of conducting these interactions. A key strategy is to encourage students to help each other: such cooperative learning processes can be less intimidating, more personal, and more dynamic than traditional approaches, while the digital content provides structure and the local staff ensures discipline.

As we conduct these experiments, we hope to turn the Digital StudyHall into a learning science testbed that allows research and practice to be intimately intertwined: experiments will be carried out in realistic settings, and the best pedagogy practices get adopted quickly and widely.

1.3 Summer Deployment

We are planning to test-deploy a prototype Digital StudyHall this summer in the Lucknow headquarters and the six affiliated village schools. The test deployment will exercise the content repository, the Postmanet-based http mechanism, the EdTV “networked thin client displays,” lecture capture and replay, and homework feedback workflow. We will work with and learn from the headquarters staff, village teachers, volunteers from the teacher training institute, and students to understand how to improve the prototype.

If the prototype turns out to be viable, we will train the staff and the volunteers to operate the system, so that the system can begin to benefit teachers and students in their daily work.

Collaborating with the staff, we will design and perform preliminary pedagogy research experiments. These experiments would explore the space of possible ways the system can be used and attempt to discover the most effective methodologies.

While the above activities will take place in the existing schools, a more ambitious goal in the longer run is to set up schools at places where none exists today. Such a school would have at least a trained staff who operates the village school equipment and interacts directly with school children. The bulk of the teaching and learning activities will rely on asynchronous interactions with the headquarters. An intermediate step to reaching this more challenging goal is to set up digitally-powered higher-grade classes in some of the existing village schools. The vast majority of the rural schools offer only classes one through five. The lack of opportunities for affordable education beyond class five is a serious problem. At places where such opportunities do exist, girls are also more likely to drop out before they reach these higher grades. During our summer trial deployment, we will set up digitally-powered higher-grade classes for girls in the premises of some of the existing schools.

Once the communication and device infrastructures are put in place of the villages, these infrastructures have the potential of supporting other types of applications and services that may improve the lives of the locals. These may include health care, communication, and commerce-related services, which we discuss near the end of this paper. During our stay, we will study the feasibility of these services in the local settings.

The overall objective of the summer deployment is to learn enough to devise a scale-up plan that encompasses the technical, pedagogical, and financial requirements of what it would take to allow more villages, more schools, and more children to benefit from the system.

1.4 Questions and Paper Organization

In the remainder of this paper, we answer the following questions:

- What are the pressing problems faced by the StudyHall schools today? How does the Digital StudyHall compare to current satellite-based approaches? We address these background questions in Section 2.
- What is the model for connecting remote learners and teachers? We address this question by discussing the learning eBay model in Section 3.
- How do we *realistically* provide connectivity between the StudyHall headquarters (which houses the centralized repository site) and the many vil-

lage schools *today*? We address this question and discuss the Postmanet and phttp mechanisms in Sections 4 and 5.

- How do we quickly bootstrap the site and populate it with good teaching content so that the site may immediately begin to benefit the village schools? We address this question in Section 6.
- How do we address the “display problem” in a cost-effective fashion: the problem of too many kids in a rural school clustered around a single computer display and few can see well? What are the potential implications of our solution? We address these questions and discuss the EdTV system in Section 7.
- How do we digitize homework, transmit it, grade it, and return feedback to kids, in a cost-effective fashion? We address this question in Section 8.
- We know we cannot effectively teach kids just by making them watch lessons on TV. Is there something fundamentally different here? How do we provide personalized interaction? We discuss some of the pedagogical research questions in Section 9.
- Is access to electricity a problem? How do we envision the Digital StudyHall to be funded in a sustainable fashion in the long run? We briefly discuss our options in Section 10.
- What are the implications of the infrastructure that we are putting in place? What other plausible synergistic applications and services can be provided with the same infrastructure? One of the motivations for examining some of the potential for-profit applications is the hope that such applications could collect revenue to fund the non-profit education and other services. We address these questions in Section 11, before we conclude in Section 12.

2 Background

2.1 Satellite-Based Approaches

It is useful to compare our approach 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 of the important advantages of our approach is that it allows high-bandwidth and any-to-any communication, which in turn enables a high degree of content customization and rich two-way exchanges, crucial for applications such as homework collection and feedback. Indeed, this theme is not limited to the networking level—it permeates every aspect of our system. For example, the content and schedules of EdTV “programs” are village-, teacher-, and student-specific. This level of customization is something that regular broadcast TVs (or satellite TVs) cannot hope to match. The virtual “market place” of education content and services provided by the learning eBay allows many more other customized interaction opportunities.

Having noted the differences between our system and satellite-based approaches, and the fact that the key components of our system share a great deal of synergy, however, we should also point out that components such as EdTV and the learning eBay are to a large extent orthogonal to the underlying connectivity technologies. It is precisely these components that constitute the “whole-system” solution that we need, the kind of components that EDUSAT lacks today.

2.2 StudyHall Background

Dr. Urvashi Sahni received her PhD from the Graduate School of Education at University of California Berkeley in 1994. Dr. Sahni has been widely recognized for her efforts to reform education in India and improve education opportunities for girls [18]. Her accomplishments include: initiating and managing a school reform project involving 62 schools, 16,000 students, and 258 teachers in rural areas of Uttar Pradesh; launching an innovative in-service program for the United Nations Children’s Fund (UNICEF) with 30,000 kindergarten and first-grade teachers in 28 districts within Uttar Pradesh; working as the director of an action research project to bolster girls’ education and serving on the state government’s Girl Child Mission to promote girls’ education; running the highly regarded StudyHall school in Lucknow, as well as operating its affiliated girls’ school and subsidized schools in rural areas.

The StudyHall headquarters school caters to students from middle-income families. Regular StudyHall classes run in the morning. The staff at the headquarters is well trained. The school has earned a reputation for excellence. In the afternoon, the school premises are used for an after-school program that Dr. Sahni runs and co-funds. This program targets girls from the urban slums, who would otherwise not have an opportunity to receive a formal education. The girls are required to do various household chores and work in the mornings and evenings, which is why the after-school program takes place in the afternoon hours. In addition to providing a precious education opportunity, perhaps more importantly, the after-school program has achieved a degree of attitudinal change: the girls have grown more confident about themselves.

The StudyHall headquarters school is also affiliated with six schools in rural villages, where the conditions are much poorer than that at the Lucknow headquarters. Each school has about 250 students. Low teacher-student ratio is a severe problem: each school is staffed by two to six teachers, who are usually not fully qualified to teach the subjects that they are required to teach. The subjects that the teachers feel least comfortable with are English, mathematics, and science. The village schools offer classes from class one to five, with the exception of one that offers up to class eight, while the headquarters offers classes from pre-nursery through pre-college (class twelve). On average, each person of the rural families (and the urban slum families which StudyHall targets) lives on less than \$2 per day.

In the summer of 2000, Dr. Sahni started a three-year pilot program exploring the use of computer technology in her schools. She put together a team comprising two government school primary teachers, two recent graduates from the teacher training institute, one ninth grade student, and two Flash software developers. This team produced science lesson units in Hindi for classes five through eight. These computer-based lesson units became integrated into the curriculum, and teachers began to use the courseware as a lecturing tool. The shortage of village school teachers trained in sciences meant that these lesson units were badly needed. The lesson units also facilitated small-team learning: under teacher supervision, students in small groups helped one another to understand the material presented by the courseware, and senior students made use of the courseware to help their juniors learn. The pilot program

also helped students gain computer literacy, which boosted students' interest in schools, and improved parents' enthusiasm of sending their children to these schools.

The preliminary StudyHall experiences have shown a great deal of potential. There are, however, a number of problems that we must address to take StudyHall to the "next level." The individual schools are not parts of a connected and cohesive whole system. Lack of physical network connectivity is only part of the problem. There is not a *system* that allows the headquarters and the village schools, and the village schools among themselves, to effectively communicate and share content. This deficiency prevents the village schools from fully exploiting the best human and content resources in the Lucknow urban environment to address their own shortages. The production of lesson units, such as the Flash content, is time- and labor-intensive. We need a faster way of making the high-quality lecturing at the headquarters available to the poorly staffed village schools. Students at the back of village classrooms complained that they could not see the computer display at the front. We need a cost-effective way of addressing such practical issues. And as we explore the use of technology in such an environment, sound pedagogical methodologies must be devised to keep up with these developments.

We would like to build on Dr. Sahni's experience in working with teachers and children, particularly those in rural areas and urban slums. Our initial steps aim at knitting the Lucknow headquarters and the six village schools into a cohesive whole so that the village schools can better benefit from the teaching resources in the urban environment. Over time, we would like to scale up the system so that more and more of the children in rural areas and urban slums can be brought into the system, enjoying a level of education that they could not have gotten otherwise.

3 The Learning eBay

At the center of what we want to build is a site that connects learners and teaching staff across time and space (Figure 1). Volunteers (who may choose to work for free) and professionals (who teach on the system for a fee) "plug" themselves into the system to play various roles. Some may develop teaching materials and make them available on the site. Some may use other people's materials and conduct teaching sessions, either virtually across the network or

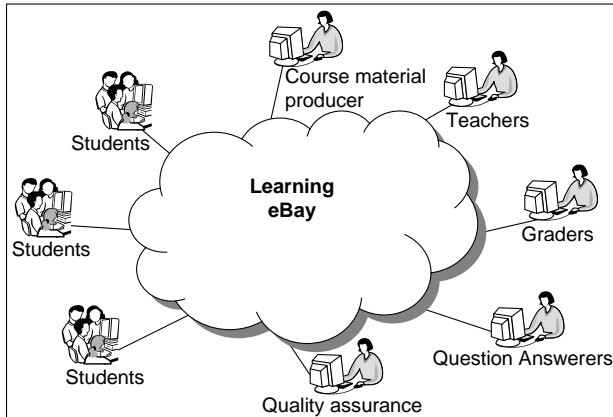


Figure 1: A “learning eBay.”

locally. Some may grade homework assignments that students have submitted to the site; the graded results are available for students to download. Some may conduct virtual “office hours.” Students and staff with limited experience or resources may tap into this online site to enrich or start their local schools. It’s also a way for teachers to get trained.

We dub this site a “learning eBay,” in the sense that this is a virtual meeting place and a “market place” that aggregates and matches “supplies” and “demands,” supplies and demands for educational resources and services.

Such a site would allow teaching professionals and volunteers to make flexible time and location commitments: for example, a volunteer can decide to spend perhaps only four hours a week grading some homework, potentially from his home overseas. This arrangement may allow us to address the difficult issue of attracting and retaining well trained and qualified teaching staff in remote regions. It may also allow us to build an education system that is analogous to an “open source” model, in which legions of content developers all over the world constantly pool their contributions towards a single coherent repository that is accessible to all. It may also allow parts of the operations of a traditional school to be operated based on an “outsourcing” model: for example, homework grading from villages can be standardized and “outsourced” to a centralized remote location (at Study-Hall in Lucknow, for example). Such an outsourcing model can address staff shortage in remote areas, ensure uniform and high standards of the outsourced operations, increase specialization and efficiency.

4 The Postmanet

To connect the headquarters repository site to the village schools, we need a connectivity option that works *today*, is relatively inexpensive, and allows us to relatively quickly scale up the number of village schools. We have examined several options. Due to the difficulty of providing commercially viable cell coverage in sparsely populated rural areas, the tele-density of cell phone usage in India as of 2003 is only 2.5% [9]. Directional 802.11 [2], while promising, is still largely experimental, and setup costs such as that needed for the erection of tall towers (on which parabolic antennas can be mounted) are very high. We are planning to use packet radio (ham radio) to provide instantaneous but low-bandwidth connectivity between a village school and the Lucknow headquarters. To transmit the multimedia educational content in both directions, however, we also need a high-bandwidth connectivity option.

4.1 What Is the Postmanet?

Making high-bandwidth wide-area Internet access pervasively available to a large rural audience is a daunting challenge. Instead of waiting for the uncertain takeoff of a number of existing and proposed technologies, which can be many years away, in a recent SIGCOMM position paper [20], we propose to turn the *existing* world-wide postal systems into a generic digital communication mechanism as digital storage media is transported through the postal “network.” The proposed system is dubbed the *Postmanet*.

While specialized solutions (such as those employed by AOL, Netflix, and some researchers working on astronomy data [7]) have emerged, they lack two key desired properties: *generality* and *transparency*: a general Postmanet should be able to cater to a variety of applications; and a transparent Postmanet should minimize manual handling of the storage media being transported. One way of better understanding the importance of these goals is to consider an imaginary Postmanet router device (illustrated in Figure 2).

A Postmanet router (or a *P-router*) is similar to a home DSL router. Instead of always forcing outgoing data through a weak wide-area network, however, the P-router writes some of the outgoing data to a mobile storage media (such as a DVD). The types of storage media used may include read-only or read-write DVDs, flash memory cards, or hard disks. We shall




Figure 2: A Postmanet router.

generally refer to these storage devices as *P-disks*. An outgoing P-disk, after being ejected from a P-router, is picked up by a postman for delivery via the postal system. The postman may also drop off an incoming P-disk, whose data appears on a user computer as if it had arrived from a conventional WAN. Therefore, unlike specialized solutions such as those employed by AOL and Netflix, the Postmanet should provide generic two-way communication, just as conventional networks do.

The user of a P-router need not manually inspect or process the content of a P-disk; the user need not manually stage or copy data; and the user need not worry about issues such as potential loss or damage of P-disks in the postal system. Unlike an AOL or Netflix user, who must know what to do manually with these application-specific disks, a Postmanet user’s only direct manual interaction with the P-router is limited to the insertion/removal of P-disks into/from P-routers. This is analogous to the fact that low-level details such as packets and routers are minimally visible to a conventional network user.

When a P-router user needs to send to multiple receivers, or when multiple applications need to share the Postmanet, ideally, it would be desirable if only a single outgoing P-disk needs to be sent per postman visit. Similarly, if a P-router user needs to receive from multiple senders, it would be desirable if there is only a single incoming P-disk that contains all the incoming data. This is in contrast to ad hoc application-specific solutions, which never allow, for example, AOL and Netflix data to be placed on a single disk. The sharing of a single Postmanet infrastructure by multiple applications and multiple users is consistent with the multiplexing and de-multiplexing jobs performed by conventional networks.

The provision of an application-neutral Postmanet “public transit” system that is easily and cheaply exploitable by any potential communicating parties is important. This is analogous to the fact that the ex-

isting Internet is such a generic infrastructure. Without it, a potential innovator who is interested in developing a Netflix-like application may need to reinvent the whole infrastructure from scratch. The co-existence of multiple Netflix-like infrastructures can lead to various forms of inefficiency. Smaller players may not be able to afford to put up their own infrastructure at all.

4.2 The Postmanet Advantages

Compared to more conventional wide-area connectivity technologies, the Postmanet enjoys several important advantages.

- *Wide reach.* The postal system is a truly global “network” that reaches a far greater percentage of the world’s human population. To leverage the postal system for digital communication, one needs no significant new investment in exotic equipment. In India, for example, facilities for daily delivery and clearance of mail exist in every village in the country [15]. The Postmanet router could be a shared resource that is deployed in locations such village schools. We also discuss ways of bringing services to individual households in later sections (using the EdTV approaches).

- *Great bandwidth potential.* While the bandwidth potential of a “sneaker net” is well known, some may consider it to be a temporary fluke stemming from the relatively poor capacity of today’s Internet. We, however, believe that this is not necessarily the case, if we examine some fundamental technology trends. Storage density of flash memory and magnetic disks has been increasing at the annual rate between 60% and 100%, and it is likely to continue in the foreseeable future. This tremendous rate of improvement is likely to be almost directly translatable to the amount of bytes transportable by the postal system for a fixed cost or in a fixed volume. Besides flash memory and hard disks, the next generation Blu-Ray DVDs can hold up to 27 GB per single-layer disc today (and 54 GB per dual-layer disc). Sony is planning to commercialize a 4-layer 100 GB version in 2007, and it has already successfully developed an 8-layer 200 GB version [13]. Commercially realizable technologies that would produce 4-layer 1 TB discs are expected to be available around 2010 [19]. One can also ship multiple units of these storage devices. As better storage devices become available, they can be instantaneously and incrementally translated into Postmanet bandwidth improvements.

In contrast, the wide-area network bandwidth growth is constrained by labor-intensive and costly factors such as how quickly we can dig ditches to bury fibers in the ground, how quickly we can furnish last-mile wiring to homes (an endeavor that can be prohibitively expensive), how quickly we can launch satellites, or how quickly we can erect WiMax (the longer-distance versions of WiFi) towers. These factors are unlikely to improve faster than the exponential growth rate of storage density. Satellite- and WiMax-based solutions may face aggregate bandwidth limitations. And the future of some of these alternatives (such as WiMax) is far from certain. Far from being a temporary fluke, the bandwidth gap between Postmanet and more conventional alternatives is likely here to stay and, indeed, widen. We do not, however, necessarily view the Postmanet as a competitor to these other alternatives. Before better alternatives become a widely deployed reality, exploring the Postmanet, an alternative that can already deliver practically infinite bandwidth today, may foster the development of and demand for sophisticated bandwidth-intensive applications, which may one day readily migrate onto alternative connectivity technologies. The Postmanet can also coexist with other connectivity technologies to complement their bandwidth limitations.

- *Low cost.* The goal of providing citizens with affordable access to postal service is typically an integral part of most nations’ postal system charters.
- *Ease of incremental adoption.* A single pair of Postmanet users can already derive useful value from the system, without having to wait for a massive-scale user community or world-wide infrastructure to develop. From this modest start, the system can grow gradually. This incremental deployment may circumvent the classic “chicken-and-egg” problem associated with the difficulty of simultaneously developing infrastructures, applications, and user populations.

4.3 Some Implementation Considerations

The Postmanet has long (but reasonably predictable) latencies. We call such a channel a High Latency High Bandwidth (HLHB) channel. Correspondingly, we call a traditional Internet connection a Low Latency Low Bandwidth (LLLB) channel. For places that have access to both an HLHB channel and an

LLLB channel (such as that enabled by packet radio), an interesting problem is how to exploit an integrated and simultaneous use of *both* channels to get the best of both worlds. For example, small requests, acknowledgements, “NAKs,” and control messages may be sent along the LLLB Internet, while large messages are staged on mobile storage devices for transmission by the HLHB postal system.

There are several alternatives that we can use to “route” P-disks from senders to receivers [20]. In a simple but effective method, well suited for our initial StudyHall environment, an end user always sends/receives P-disks directly to/from a single data distribution center (called a *P-center*). (The P-center would be co-located with the StudyHall headquarters in Lucknow.) Although any centralized solutions have obvious disadvantages, an important advantage of this approach is that each end user handles only a single P-disk, regardless how many other sites he communicates with per postman visit: as the P-center copies data from its incoming P-disks to its outgoing P-disks, it first demultiplexes incoming data and then re-multiplexes outgoing data, minimizing the number of P-disks handled in both directions. Inexpensive robotic arm-operated, multi-drive DVD writers that can process about 600 DVDs per day already exist today and they can keep manual labor cost to a minimum.

Among the many potential applications of Postmanet that we have discussed in our position paper [20], a most promising one is the delivery of basic education to impoverished rural areas: the wide reach, huge bandwidth, and low cost advantages of the Postmanet makes it an attractive mechanism for transmitting multimedia content between StudyHall and the village schools in both directions.

5 phttp: Postmanet-based http

5.1 What Is phttp?

The central repository housed at the StudyHall headquarters in Lucknow is accessible via the conventional Internet so teaching staff and volunteers who have “normal” web access from either within India or from overseas can contribute to it. To make the repository available via a web interface to isolated village schools, however, requires an unconventional networking approach, and one of the most important applications of the Postmanet discussed in the previous section is to deliver web access. Instead of trans-

mitting web request and reply messages across a wire, in this section, we discuss a system that automatically encapsulates these messages entirely on DVD media transported by the postal system. We call it “Postmanet-based http,” or *phhttp*.

Under *phhttp*, a staff member operating a village computer can request (pull) content from the repository, and the headquarters staff can push content to the village schools. They can also perform other operations such as search and browse. Note that to simulate a web experience, we need to handle not only static data, but also actions that are normally effected by scripts and programs executing on web servers. Under *phhttp*, almost all of the DVD-related activities are automated so the users are minimally aware of them.

The operations of *phhttp* are related to, but also very different from, those of “offline web browsing.” An offline web browser *eventually* gets directly connected to the Internet, while a village computer may never do. Offline browsers can deliver only minimum functionality when server-based scripts are involved in satisfying requests, while under *phhttp*, we explicitly migrate some code (as well as data) from the server to the clients so the client browsers continue to function in absence of any conventional connectivity.

We illustrate the operations of *phhttp* using the examples in Figure 3. In panel (a), the repository “publishes” an initial bootstrap DVD destined for each village school. Each bootstrap DVD contains enough initial data and a collection of scripts that would allow a client to interact with the system without any conventional connectivity. Note that the total amount of the data contained in the repository is far larger than that can be fit on a DVD so the purpose of the bootstrap DVD is just that: bootstrapping. We use a robotic arm-operated DVD processor to mass-produce these DVDs so the amount of manual intervention is minimum. In this initial step, all the bootstrapping DVDs may contain identical content; but as we shall see in later steps, in general, the outgoing DVDs produced by the repository site contain different content.

In panel (b), a bootstrap DVD arrives at a village, and is inserted into a Postmanet router. Its content is automatically copied onto the local disk. Upon completion of the copying, the rewritable DVD media is automatically erased. Again, no manual intervention beyond inserting the DVD is required.

In panel (c), a staff member at the village interacts

with the computer at a time of her choosing. Enough data and scripts from the repository, transported by the incoming DVD, has been copied to the local hard disk to allow the staff to interact with a subset of the full repository site. In our case, for example, an up-to-date version of the full catalog (or “metadata”) of the repository content is always copied to the village school hard disks to enable browsing and searching. Scripts that allow the staff to place “orders” of chosen repository content, as well as those that allow the staff to upload selected local content (such as digitized homework submissions) are also deposited onto the local hard disk. As the staff interacts with this local view of the site, she may generate various requests, including orders (or download requests) of repository content and upload requests of local content (such as homework submissions). These requests are buffered on the local hard disk.

In panel (d), shortly before the the next postman visit, the buffered requests on the local hard disk is automatically burned onto a DVD (which is likely to be the same rewritable DVD that is erased in panel (b)). Again, the user is not involved in activities such as file copying: all that is required is removing the DVD from the Postmanet box and handing it to the postal worker.

In panel (e), all the incoming DVDs from the many village schools are collected at the StudyHall and they form an incoming DVD stack. The robotic arm-operated DVD processor automatically reads the content of each incoming DVD onto a hard disk, and erases each incoming DVD, forming a stack of blank DVDs. Next, the repository machine processes the requests from each village school (that are now stored on a local hard disk) and attempts to satisfy them. If the request is for uploading content into the repository, the data is copied into the repository database. If the request is for downloading content from the repository, data is copied out of the repository database and is placed in an outgoing DVD image on the local hard disk. At the end of the processing, an outgoing DVD image is generated for each village school on a local hard disk.

Now the cycle repeats and we are back to panel (a)—we are ready to generate another stack of outgoing DVDs, using the stack of blank DVDs produced by the erasure step in panel (e), and the outgoing DVD images placed on a hard disk (also in panel (e)). Unlike the outgoing bootstrapping DVDs produced initially, however, the new stack of outgoing DVDs generated now can contain different content

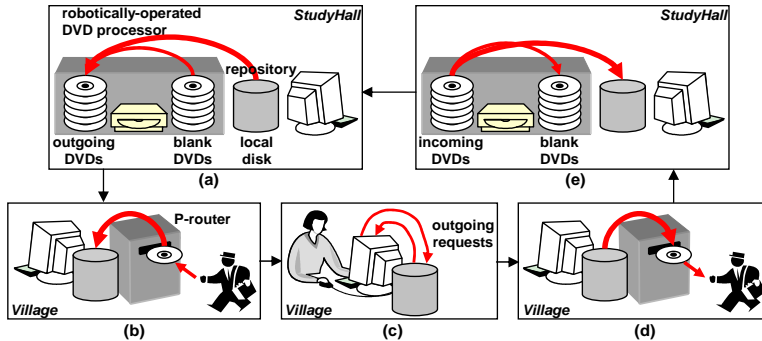


Figure 3: An example cycle of phttp operations. A key aspect of this process is that there is minimal manual involvement in the DVD-related operations.

(to satisfy different requests from the village schools). And when these DVDs reach the schools in panel (b), the requested content will be automatically copied onto the village local hard disks.

5.2 Observations

While the phttp example steps that we have examined in Figure 3 are intuitive, and a savvy computer user could have done some of this manually, we emphasize that one of the important aspects of our approach is *transparency*: almost all of these steps are automated and not much of the DVD-related operations is visible to end users. This automation and transparency is important for at least three reasons: (1) if our goal is to scale up the digital StudyHall to encompass a large number of schools and students, manual intervention becomes tedious and infeasible; (2) automatic handling of exceptional events, such as lost or damaged DVDs, is also important if we need to scale up; (3) figuring out what scripts to place on DVDs for any particular site is not a trivial task that ordinary users or even site administrators can figure out for themselves—we need a well-defined phttp API that “splits” centralized server scripts into pieces that are distributed at the client and server sites to intelligently support asynchronous operations. In addition to their large capacity, low cost, and small weight, the reasons that have compelled us to choose DVD media include the fact that we can use commercially available (and cheap) robotic arm-operated DVD processors, a crucial part of our automation strategy.

Although our discussion of Figure 3 assumes that all communication is carried out by transmitting DVDs, as we have discussed in Section 4, a low-bandwidth packet radio link, if available, for example,

can be used to complement the high-bandwidth DVD transmission. Small data items, such as the catalog of metadata, can be made available on a packet radio connection, so are small request messages and notifications such as receipt acknowledgements, NAKs generated as a result of damaged DVDs, and retransmission requests.

The discussion in this section has focused on the StudyHall repository application. The phttp approach can be used to support other applications. For example, a user may receive a large digital catalog of Amazon.com on an incoming DVD, browse the catalog and place purchase orders that are transmitted back via an outgoing DVD. Or a user may send a Google search request and receive a reply DVD that not only contains the search result page, but also a crawl of the top-ranked sites. Or we may build a voice/video mail application. An attraction of a general Postmanet or phttp infrastructure is that all such applications may share the same infrastructure and, indeed, the same incoming and outgoing DVDs.

6 Content Capture

Our aim is to quickly populate enough content in the repository so that the village schools consuming the digital feed from the StudyHall repository may immediately start to benefit from it. A critical mass of initial content may also provide a good starting point that attracts participants to attempt to improve the content, with the knowledge that such improvements may immediately see practical use. Otherwise, we might risk the chicken-and-egg problem of having to simultaneously (and slowly) develop content and attract participants. Carefully authored content such as Flash movies can be very time- and labor-intensive.

We cannot rely on such content alone as a means of quickly populating the repository.

Our approach is to record the live lectures at the StudyHall headquarters in Lucknow. The Lucknow school is taught by a better trained professional staff. Our goal is to allow rural children to receive a level of education that is comparable to that enjoyed by middle-class children in urban areas. Over time, we expect participants from other places would start to contribute toward the repository with better and better content. Of course, we recognize that merely replaying live lectures in front of rural kids, by itself, is not sufficient. We will address the issues of homework and interactions in later sections. Nevertheless, capturing and replaying high-quality lectures is still an important part of a bigger solution.

We will experiment with several different means of digital capturing. (1) Real time hardware MPEG4 encoders are available (and cheap). The captured DVD-quality videos are sufficiently compact and hard disks have become sufficiently big and cheap that we can realistically capture and permanently store a very large number of lectures and other types of interactions in the repository. (2) For content that is shown on computer displays at the Lucknow school, the evolving display content (along with sound in the environment) can be directly captured on the machine driving the display (using software such as CamStudio). Such content can be potentially captured at higher resolutions, lower frame rates, and/or smaller file sizes than those of conventional DVD-quality video. (3) We have also built a system that utilizes a still digital camera: the camera would be programmed to release its shutter at a constant rate, and the resulting frames, along with a recorded sound track, can be “stitched” together to form a video, a video that potentially has a very high resolution and a low frame rate. This can be accomplished using tools such as those based on the “Synchronized Multimedia Integration Language” (or SMIL). (4) We will also store many other types of digital content in the repository. These include digitized homework and feedback, question and answer sessions, etc. We will discuss these issues further in later sections.

Initially, we plan to perform passive capturing, with minimum intrusion on the live sessions. In the future, we may experiment with various degree of teaching staff cooperation, both during and after the capture sessions. We will also work with volunteers from a neighboring teacher training institute. These volunteers can perform various forms of post-

processing to improve the captured content, stage their own lectures to be captured, produce feedback sessions for homework submitted from village schools, and run asynchronous digital question-answer sessions. In the process, these volunteers gain exposure to technology and real-life training interacting with village students (digitally).

We expect content captured elsewhere will be contributed to the repository as well. In the long run, the repository could develop into a peer-to-peer architecture if there are multiple “centers” of content accumulation.

The cost of the equipment needed for the capture activities can be kept very low. Furthermore, although it is critical that we do diligent “penny-pinching” to keep the cost of the per village equipment low (because such cost will be scaled up many-fold as we increase the number of schools covered by the Digital StudyHall), the extra cost incurred on slightly more expensive capture equipment at a small number of centers of teaching excellence should not matter.

7 EdTV: Thin Client Displays

7.1 What Is EdTV?

One problem that we need to solve with a digital school is the “display problem:” the problem of too many kids clustered in front of a single computer display and few can see well. (Each of the village schools affiliated with StudyHall today, for example, serves about 250 students, divided into a few classes.) The obvious solutions have practical problems. Increasing the number of displays is costly. Projectors are both too expensive and too power-hungry. In order to scale up the Digital StudyHall to eventually encompass a large number of schools and students, it is critical that we use only cost-effective solutions in the village schools.

Our solution is (partially) shown in Figure 4. A graphics card capable of outputting RCA video signal is installed in the village computer. The analog video signal is sent to a small TV transmitter. A number of regular TV sets are placed in the classroom, each of which receives the air signal coming from the transmitter, and displays the village computer display content on the TV screen. Only a small number of students need to huddle in front of each TV screen, which essentially acts as a computer display.

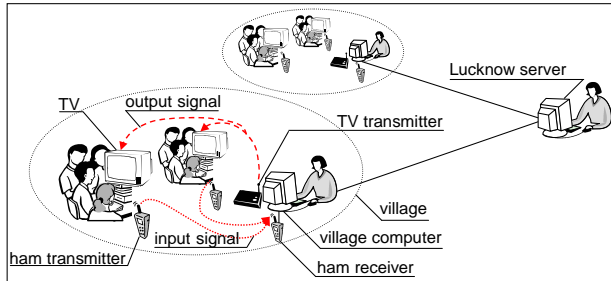


Figure 4: EdTV.

This solution is inexpensive. The graphics card costs \$12. The TV signal transmitter costs \$37. India has a very healthy used TV market and TV sets can be had cheaply too. This solution is also energy-friendly. Twelve-inch sets consume as little as 20 W, far lower than that consumed by almost all laptops, and a projector can consume ten times or more.

Under this arrangement, the TV sets do not need to be inside the village classroom to work: TV sets anywhere in the village within the transmission range of our transmitter can receive the display signal as well. Energy permitting, we can broadcast beyond the classroom far past regular school hours, so for example, a kid who misses a school day due to work during the day can schedule time slots on the transmitter at night to catch up. (A recent survey of some of the StudyHall village schools, for example, shows attendance levels of 50% during mango-picking season.) There are many other ways we can use the system, including homework feedback sessions and other potential applications. We can also introduce “input” devices that allow viewers to interact with the system. We will discuss these issues in later sections. We call these local TV “networks” *EdTV*.

7.2 EdTV Is *Not* TV

It is worth noting that despite its initial promise as an education tool, TVs have not had the best reputation serving educational purposes. We would like to point out that EdTV is *not* TV. We say this for two types of reasons. First, from a pedagogy standpoint, we will experiment with various types of personal interactions that we discuss in later sections. Second, from a device architecture standpoint, the EdTV screens are more similar to “thin client displays” than regular TVs. We now discuss this distinction.

Regular TVs are “mass media devices”—a very small number of content producers dictate what and

when a large number of content consumers watch. Computer displays are exactly the opposite—they are “personal media devices:” they allow many-to-many communication, which can carry highly customized content and can occur at flexible times of the communicators’ choosing. EdTV may not be as “personal” as regular computer displays that we know, and it occupies an interesting compromise point in between, but it is probably in spirit much closer to the “personal media device” end, allowing, for example, village-, teacher-, and student-specific content customizations that “mass media devices” are utterly incapable of. The homework feedback sessions that we discuss in the next section are an example of such customizations.

7.3 EdTV Is Not WebTV

In some sense, the EdTV approach is similar to that taken by WebTV: both attempt to leverage an existing large install base of legacy devices and an appliance usage metaphor that people are familiar and comfortable with to deliver new applications and new services. EdTV is not as personal as WebTV, and it is not meant to support the existing web experience directly, but it has some important advantages compared to WebTV if we choose and design our applications carefully. EdTV does not require a dedicated per TV conventional network connection. The hardware requirements at the consumer end are minimal (and we discuss cheap “input” alternatives later). EdTV has no setup requirements at an end users’ end and the usage metaphor would be simpler. And if the village computer powering the EdTV transmitter is backed by a Postmanet connection (either wholly or partially, complemented by a low-bandwidth modem), the shared “backend network” would be bandwidth-rich, cheap, and pervasively available today. And these backend network advantages are passed onto the end EdTV users.

Past experiences with Information and Communication Technology (ICT) projects in developing countries have concluded that finding ways of “aggregating demand” is an important strategy of bringing ICT services to low-income people in a cost-effective fashion [14]. Another lesson is that innovative ways of combining “old media” with new technologies can be effective. One example is the Sri Lankan Kothmale Community Radio, where villagers use cell phones to call a live radio station operator, who would look up the requested information on the web in real time,

and broadcast verbally the answers over the radio wave to a local community.

The EdTV approach is consistent with these lessons. What is shown in Figure 4 is a three-tiered system. The bottom tier, the TV sets (with possible addition of some input means that we discuss later), consists of cheap, specialized and simple devices, combining metaphors of both old and new media. The shared middle tier, the village computers, is somewhat more expensive, more general, and more capable. An important function of this middle tier is to aggregate demand. The data temporarily stored in this layer is potentially expendable. The top tier, the Lucknow server, is more professionally managed, and contains the reliable and ultimate “truth” in its repository.

It is also interesting to compare the EdTV and conventional kiosk approaches, both of which allow end users to share relatively expensive computer and connectivity resources. EdTV has two significant advantages compared to kiosks. First, EdTV allows one to access services from the convenience of one’s home at a time of one’s choosing. One may still get a “busy signal” but that is a price of sharing. Second, EdTV allows multiple people in front of multiple TV screens to share a collective experience. Such an experience can be more compelling if the application is specifically designed to support this mode of sharing. One such example is a specially designed auction program. (See Section 11.3.) Lack of privacy is a potential issue with the single-user EdTV scenarios, but the multi-user scenarios in fact exploit it to provide unique and interesting sharing experiences.

7.4 Extensions

We now discuss several extensions to the basic EdTV setup. The transmitters that we are looking at have ranges between 500 feet and several miles. To extend the range when necessary, we may build cheap signal repeaters of our own, by gluing additional transmitters with used VCRs. (Only the TV receiver parts of the VCRs are needed for our purpose.)

Another useful extension is providing the end TV watchers with means of sending input to the village computer. As shown in Figure 4, an inexpensive custom-made ham radio transmitter acts as a “TV remote,” whose commands are picked up by a ham radio receiver, which is in turn connected to the village computer. (The same ham radio receiver co-located with the village computer can handle both

communication with the Lucknow server and commands from the EdTV input devices.) We may also use walkie talkies: a receiver walkie talkie “speaking” into the microphone at a village computer running voice command recognition allows end users to use their sender walkie talkies to effect some control. Carefully designed user interfaces would be helpful, especially in shared viewing scenarios. One way of looking at EdTV is that the output TV signals and the input radio signals are a way of truly bridging the “last mile.” However, we should not overlook “low-tech” solutions. For example, a human operator may replace the role played by the voice recognition software. Or more simply, a village town hall meeting can be the setting where the schedule of the local EdTV is determined.

And instead of TV, we can apply a similar approach to radio signal transmission. Certain types of curriculum, such as the teaching of spoken English, can benefit from an audio-only approach. Songs sung by local kids in schools could be attractive radio content: past experiences suggest that such elective art-related programs, when introduced to otherwise dull daily routines in primary schools, can play a significant role raising interest and satisfaction levels in schools.

In later sections, we discuss how we perform customized educational interactions, such as providing homework feedback, over EdTV. We also speculate about other non-educational applications that one may run over EdTV.

8 Homework Feedback

8.1 The Digital Homework Workflow

Our aim is to scale up the Digital StudyHall so that it provides digital feeds to many rural schools that have minimum experienced staff resources. Such a school may be staffed by a computer operator who receives some training from the StudyHall but may have otherwise little domain-specific expertise in the subjects taught at the school. Obviously, playing the lectures recorded elsewhere is not sufficient. Handling homework feedback well is important if we were to make such schools work. We discuss other forms of interaction in the next section.

An important principle we need to observe in our homework solution is, again, low cost, a necessary requirement if we were to scale up the number of these schools. Due to the high cost of ink cartridges

and printer paper, for example, we are shying away from solutions that require a large amount of printing. Instead, we will experiment with the following digital “workflow.” As we discuss later, this process is designed not only to provide to rural children a comparable homework experience as those received by their urban middle-class peers, but also in fact to potentially *exceed* the quality level of the “normal” homework feedback process that we are familiar with today, while also keeping the cost down.

- *Digitizing.* School children still do homework the way they used to: using pencils and paper. (Lower-quality “non-computer” papers are cheap.) Each school is equipped with a digital camera, which either the staff or the students use to digitize the written homework. (We choose to use digital cameras instead of scanners due to the cameras’ speed, versatility, portability, and simpler power requirements.) Each school is also equipped with a microphone, which can be used to digitize voice. A webcam can be used to record video to add a more personal touch in certain types of interactions and generally enhance interest levels, but it is not strictly necessary. All this accessory equipment should be inexpensive.

- *Transmission.* The digitized homework is “uploaded” into the repository by the village staff using the phttp mechanism described in Section 5. Phttp is in some sense an ideal transport for our purpose because (1) homework submission and feedback are intrinsically asynchronous processes, and (2) the tremendous bandwidth potential of phttp allows us a great deal of flexibility in terms of the type and content of homework and feedback that we can transmit: we are free to use all manners of multimedia content that can provide good end user experiences (experiences that kids are either already familiar with or experiences that are engaging) without worrying about over-taxing communication links or incurring high per-byte cost; one is unlikely to match these advantages using other communication means.

- *Grading.* As part of the phttp mechanism, the robotic arm-operated DVD processor at the Lucknow headquarters automatically uploads the incoming homework submissions into the repository. The grader staff (or volunteers elsewhere) downloads the submission at a time of their choosing. A grader loads the homework into a batch image management software and uses a tablet pen to scribble on the homework submission.

In addition to this “written” feedback, a grader also prepares a video feedback, which is produced by a screen capture process (using software such as CamStudio), as the grader scrolls through the graded homework and speaks into the microphone to provide a verbal soundtrack. A teaching assistant of an introductory Princeton CS course has been using this process to grade homework, and we have found the process to be quite natural to use.

In the context of the Digital StudyHall, we plan to produce two types of homework feedback: *collective* and *individual*. Collective feedback is a video that a whole class views together, and is designed to cover issues such as common mistakes. Individual feedback provides more personalized help, especially for the kids who appear to be having trouble. The distinction between these two types of feedback allows us to reduce the total length of the feedback sessions.

The hardware needs of a grader is also minimum: an inexpensive tablet pen and a microphone are all that is needed. (The video is produced via screen capture so there is no camera involved.) The modest equipment needs make it easier to scale up the number of graders in urban centers such as the Lucknow StudyHall headquarters, and also make it easier for individual volunteers elsewhere to join in and help.

When done, the grader uploads the graded homework, along with extra feedback files (such as video files), back into the repository. This is likely done with the high-bandwidth local area network at the Lucknow headquarters, or a conventional wide-area network if the grader is an outside volunteer. It is also possible, however, that a grader is located at a bandwidth-starved location, so she could use the phttp mechanism to interact with the repository just like a village school staff does. In all cases, the grader instructs the repository to push (again, via phttp) the homework feedback to the village schools that submitted the homework.

- *Feedback.* When the graded homework and feedback arrives at the village school over phttp, data is automatically copied onto the local hard disk. At a time of her choosing, the staff goes over the homework feedback with her students. The collective feedback video may be played over EdTV to all students in the classroom. The individual feedback video for a student may be scheduled to be played at a time of the student’s convenience, either in front of the computer display, or on an EdTV screen at or near the student’s home after regular school hours. The stu-

dent may borrow a walkie talkie as an input device, which allows her to pause, rewind, zoom in, or perform other operations to the content being delivered over EdTV.

If necessary, the staff or the student may choose to review some raw images of the graded homework. If the number of students is large, we may increase the number of EdTV channels by equipping the village computer with additional RCA signal-capable graphics cards and transmitters, at the cost of about \$50 per additional channel. Homework feedback is an example application of EdTV that showcases our contention that EdTV is *not* TV: it allows customized content and customized control that is impossible with “mass media devices.”

Although we anticipate using EdTV to provide a large fraction of the feedback, we do not necessarily rule out other means. Printing out *some* content, as long as the scale of printing is not massive, is perfectly fine. “Answer books,” printed on cheaper papers, like textbooks, could be shipped from the StudyHall headquarters in advance, kept by the village staff, and provided to students when needed. This combination of the low-tech and the high-tech is a theme of the approach that we are taking to maximize communication effectiveness while minimizing cost.

8.2 Observations and Implications

Having described the remote homework feedback process, we now make some additional observations about the process. Despite the fact that the rural villages may not be staffed by teachers or graders with much domain-specific expertise, the homework feedback experience received by students could in fact be potentially better than that received by middle-class students today: the video feedback sessions, for example, can be easily and cheaply made and provide a much richer and personal learning experience.

The process that we have described is useful not only for doing homework, but also useful for other interactions in general, such as less structured question and answer sessions. (Of course, this is not as good as conventional real-time video conference sessions; the asynchronous interaction may nevertheless still prove valuable.)

All the feedback materials are permanently stored in the repository so they can be reused. Some of the reuse scenarios allow us to partially overcome the http delay problem. For example, a local staff

could play a good previously recorded feedback session (potentially made for other students elsewhere) almost immediately after the homework submission of the current class is made. More personalized feedback for the current class would come later. So the feedback a student receives includes both instantaneous feedback (although the feedback is not terribly customized) and customized feedback (although the feedback is delayed). More details of how this tradeoff may work can be found in an earlier paper of ours [11]. One may also reuse previously recorded material by selectively mixing and matching snippets with newly recorded material.

Even for a school that *is* locally staffed by well-qualified graders, the repository may still prove valuable. A grader may choose to occasionally download and play some of the feedback recorded at other times or at other places to augment her own feedback sessions. The homework graded locally may be digitized and uploaded into the repository anyhow, so that this content may benefit students at other places or at other times. One of the more radical (but potentially very promising) approaches that we are considering is to have the best higher-grade students help other students, potentially located elsewhere. Such “cooperative learning” may not only give us “extra hands,” but may also prove to be a more personal and less intimidating learning experience for the students. We discuss such approaches more fully in the next section.

9 Pedagogy and Learning Science Research

Although the homework feedback mechanism discussed above allows a fairly sophisticated asynchronous interaction mechanism, it is still probably safe to say that playing the captured lectures alone will not be sufficient for educating young children. In this section, we discuss pedagogy research targeting our Digital StudyHall environment: the staff member manning a village school may have various degree of domain-specific knowledge, ranging from none to something quite extensive. The question is what we can do, by utilizing such a staff member, to augment the digital feeds to and from the Digital StudyHall headquarters to make the teaching process in a village school more effective.

9.1 The “TVI” Experiences

We draw our inspiration from the “Tutored Videotape Instruction” (TVI) program conducted at Stanford in the 1970s [6]. In the TVI paradigm, minimally-edited videos of unrehearsed lectures are viewed by small groups of students assisted by a *facilitator*. The facilitator is not an expert teacher; instead, the (simpler) job of the facilitator is to pause the video tape when questions arise during the play back of the taped lectures. Such questions can be asked by students who are currently viewing the tape, the instructor on the tape, or the students captured on the tape. During such a pause, the facilitator attempts to guide the students toward a resolution of the question using a discussion format, before he resumes playing the tape. When not many questions are asked, it is also the facilitator’s job to instigate more questions and discussions.

Over a number of years, careful studies were performed to compare the performance of the students participating in the TVI program against that of three other groups: one was the group of students in the live classrooms; the second was the group of students who viewed the lectures live on close-circuit TV and were given the option of interacting with the instructor using phone hookups; and the third was the group of students who viewed the delayed taped lectures without any augmentation. We call these three groups the “live” group, the “interactive TV” (or ITV) group, and the “TV” group, respectively.

While it was not surprising that the TVI group outperformed the TV group, what was especially interesting was that the TVI group outperformed *all* three other groups, including the live group, consistently over the years! Indeed, the TVI students appeared to lag behind the live (on-campus) students in terms of their admission qualifications; yet the TVI mechanism appeared to have more than compensated for this initial disadvantage.

The conjecture is that there is more systematic interaction built into the TVI model than even that in the live classroom, and this increased amount of discussion in a less inhibiting environment also fosters the formation of a group discovery process that builds communication and team skills, all of which ultimately contributing to an enhanced learning process [12]. Variations of the TVI approach have been tried subsequently [5, 1], and similar degrees of success are observed.

The TVI experiences point to several important

principles that are highly relevant to our Digital StudyHall. (1) Although not sufficient by itself, good captured content, when replayed, can be a key foundation upon which a digital school curriculum is based. (2) Systematically instigating interaction with and among students can effectively complement the previously recorded material and significantly enhance the effectiveness of the learning process. (3) Successful instigations of interactions can be effected by relatively simple means (such as a non-expert facilitator pressing the pause button) without necessarily requiring a sophisticated, time-consuming and expensive content authoring process. (4) Group learning can play a key role, especially in an environment where the local staff expertise is limited.

9.2 TVI and the Digital StudyHall

Compared to the TVI environment, the Digital StudyHall faces some very different challenges and opportunities. An obvious challenge is the fact that the students we have in the Digital StudyHall are young children, and unstructured discussion may not be best suited for this audience or for the types of the subject matters covered. On the other hand, the technologies available to us today have advanced far beyond the VHS tapes and the pause button: digital technologies potentially can allow far more sophisticated options than pause-and-resume. Yet, opportunities afforded by the digital technologies have yet to be fully explored.

Having said that, we recognize that the original pause-and-discuss TVI model has very important advantages: it is simple—it minimizes the amount of potential intrusion on the original lectures; it requires almost no post-processing of the recorded material; and a facilitator has a relatively simple task. While it is important to strive to preserve these advantages, we also recognize that there may be a “gradient:” while the original TVI model lies at one end of the spectrum in terms of minimizing intrusion, post-processing, and facilitator options, increasing levels of sophistication along this spectrum may still be worth exploring. We are helped in these explorations by a skilled and cooperative StudyHall headquarters teaching staff (who may not only tolerate some degree of intrusive changes of their normal teaching routine, but indeed also be eager to experiment with and embrace the new technologies), a digital repository that collects all teaching materials and encourages incremental tinkering and experimentation (so post-

processing efforts of the captured content across time and space can be harvested and preserved), and a village school local staff who may have different degrees of expertise and skill (so they can “facilitate” in ways beyond pause-and-resume).

We now hypothesize some example types of instigations of interactions that may play similar roles as that played by the pause-and-discuss format in the TVI model, but are potentially more appropriate for our target audience and course subjects. A five-minute in-class quiz, for example, can be built into each lecture conducted at the headquarters site. The subsequent live grading and feedback are captured along with the lecture. When this lecture is replayed in a village school at a later time, the same quiz is administered, and the previously captured grading and feedback session can serve the role of the discussion instigations in the original TVI model.

The homework feedback sessions, as discussed in Section 8, are also examples of instigating interactions. As discussed in Section 8.2, such feedback sessions may consist of instantaneous feedback (based on previously recorded material) or delayed feedback (based on tailor-made material for the specific students). The same homework feedback mechanism can also be used to conduct less structured question-and-answer sessions with the remote teaching staff at the Lucknow headquarters. What can make these interaction sessions more useful, in the context of the TVI model, is to carefully interleave and integrate them into the flow of lecture playbacks.

Yet another possible type of interaction that we can conduct in the Digital StudyHall is based on the recognition that the local facilitator, though possibly lacking domain-specific knowledge, is far better at certain tasks, such as pattern recognition, than any artificial intelligence program. A challenge is to design content streams that incorporate the facilitator intelligence as part of the “system.” For example, the content stream can instruct the facilitator to identify certain patterns in students’ verbal or written responses and choose a corresponding followup content snippet to play.

In all these interactions, just like in the TVI model, one of our goals is to encourage students to help students. In our teaching experiences, we have often seen some of the best-performing students emerging as “leaders,” helping their peers learn. Such cooperative learning processes can be less intimidating, more personal, and more dynamic than that with authoritative figures. An important task of the facil-

itator is to foster the formation of a positive group dynamic that is conducive to cooperative learning. Indeed, we will explore employing the best upper-classmen as facilitators or “facilitator assistants” in lower-grade classes. We will also explore encouraging student interactions across the wide area (such as those between urban students and rural students); their different perspectives may further enhance the learning experience.

While encouraging group learning is important, we also recognize that we need to provide structure and ensure sound resolutions, so in simple cases, for example, the framing questions and correct answers would be provided in the pre-recorded content stream, minimizing digressions during these group interactions. We also note that the types of interactions that we have discussed are only some examples: they are not meant to be exhaustive, and only through rigorous evaluation (which we discuss next) can we know their effectiveness.

9.3 A Learning Science Testbed

Education experts have long lamented the severe disconnect between research and practice: it is difficult to perform learning science research in realistic settings; and it is even more difficult to have research results adopted by front-line teachers working in the trenches [4]. The Digital StudyHall presents a unique and efficient testbed that allows research and practice to be intimately intertwined in an ongoing basis.

For example, a reputable researcher can design experiments for evaluating the effectiveness of the different instigation methods under the TVI model (discussed in the last subsection) by providing several different digital feeds and different instructions for the facilitators. These digital feeds will be pushed to the test schools and success metrics (such as in-class quiz scores) are gathered. Methodologies that do not work are identified and they are either discarded or further improved. Methodologies that prove successful would become more widely adopted as the corresponding digital feeds are picked up by subsequent teaching sessions. Further improvements to these superior digital feeds can be made over time, potentially by other practitioners. The content in the repository, like some on the web (to some extent), is self-selecting: the best gets better and gets used more often. In some sense, the system has memory.

Our hope is that the Digital StudyHall not only allows learning science experiments to be easily con-

ducted in realistic settings, but also “short-circuits” the research-to-practice turnaround time so that the best pedagogy ideas get adopted quickly and widely.

10 Other Issues

10.1 Electricity

The six village schools currently affiliated with StudyHall all have access to electricity. (One uses solar panels, while the rest have either intermittent or stable access to the grid.) In India, however, 40% of the households, and 60% of the rural households do not have access to electricity. Although access to electricity is not a main issue for our summer deployment, it will pose a challenge in the longer run as we attempt to bring more schools into the Digital StudyHall system.

For sites that have no access to electricity today, we have to rely on one of two solutions. One is to install a solar system, which costs \$600 or higher. The second is to rely on car batteries, which are periodically brought back to towns to have them charged. This is a solution that some villagers are already using today to power their TV sets. We will investigate a minimal hardware setup for a village school powered by car batteries so we can make the most out of them.

10.2 Funding

One question we are often asked is how we envision the Digital StudyHall will be funded in a sustainable fashion in the long run. We see four stages (or options). In the initial research and development stage, we will need support from grants and donations from NGOs and corporations. Beyond this initial stage, we speculate that we may be supported by three possible sources. The first is an “open source” model, which relies on a legion of distributed volunteers to maintain and improve the system. The second is local government support. Once we are able to demonstrate that the Digital StudyHall is a cost-effective approach to delivering basic education, and that a highly qualified teaching staff can have a bigger impact when working as a part of the system, rational government decision makers may choose to spend some of their education budget on such a system. Indeed, ideally, we would like to turn over the operation of the Digital StudyHall entirely to a local staff as soon as it matures. The third potential source is revenue collected from

for-profit applications and services “spun off” during the process of the Digital StudyHall development. These are the types of applications that can leverage the same hardware and software infrastructure. Revenue collected from users of such applications may be used to “cross-subsidize” the non-profit education component. We discuss some of these hypothetical applications in the next section.

11 Other Applications

Several key components of the Digital StudyHall infrastructure that we have described in preceding sections have implications beyond the education applications—other potentially compelling applications or services may leverage these components as well. These components include:

- *The Postmanet and the phttp communication mechanisms.* This is one of the key enabling components that allows us to have pervasive, high-bandwidth, and low-cost asynchronous connectivity to just about any place, including the remotest areas. (See Sections 4 and 5.)
- *The phttp-accessible content repository.* The content repository is a simple abstraction that allows operations such as upload, download, pull, push, browse, and search. In addition to the conventional web interface, we also make these operations available via phttp. And a robotic arm-operated DVD processor co-located with the repository site allows us to automate almost all aspects of the phttp operations. As we shall see, the content repository is a sufficiently general and powerful abstraction, upon which other more specific applications can be built. (See Section 5 for the phttp operations on the repository and Section 8 for an example of how people interact through the repository.)
- *The EdTV thin client displays.* The EdTV approach allows us to leverage a large install base of legacy devices, namely TVs, and turn them into thin client displays powered by a shared computer. This approach solves several difficult problems: it lowers the cost of end user devices, and it truly bridges the last mile by leveraging TV signals in one direction and input mechanisms such as radio control signals in the other. EdTV opens up a “portal” for interesting computer-based applications into a poor household that would have otherwise been beyond the reach of more conven-

tional technologies. (See Section 7 for the EdTV mechanisms and Section 8 for an example of how people use EdTV).

In this section, we discuss several hypothetical applications that can be built on top of some or all of the above key components. These applications fall into two categories. One includes non-profit services, such as the “health care eBay” that we discuss next, which happen to share a great deal of synergy with the “learning eBay” service that we have discussed so far, in terms of both their infrastructural needs and philosophy. The second category includes applications that may be potentially commercially viable. One of the reasons why such applications interest us is the potential of such applications’ generating revenue that can be used to fund the non-profit services, making them more sustainable.

11.1 The Health Care eBay

Efforts of improving rural health care face many similar challenges as those faced by providing good primary education: lack of well-trained professionals working and living in remote areas, a severe communication barrier due to the difficulty of providing pervasive connectivity, and the requirement for cost-effective solutions. Cost realism is especially important in a resource-limited environment where one must face the stark tradeoff between, for example, more medicine versus better communication. This was perhaps best illustrated by Bill Gates’ unequivocal declaration at one time that “poor people need medicine and not computers.”

In reality, however, this choice is more nuanced and is not necessarily a zero-sum game. The old truism, “an ounce of prevention is worth a pound of cure,” has a proven track record, as numerous experiences have shown that systematic availability of preventive care and early intervention can significantly reduce the overall cost of the health care system. No less a spokesman than the president of the National Hospital Association made a statement in a national interview that prevention is going to be the important focus of health care. The trick is to find a “sweet spot,” where a relatively inexpensive investment in an effective communication and interaction infrastructure can enable (and, indeed, encourage) the poor to proactively seek information and early care.

11.1.1 What Is the Health Care eBay?

In Section 3, we discussed the metaphor of a “learning eBay,” a site that connects teachers and students, and “supplies” and “demands” for education services across space and time. In this subsection, we discuss a similar metaphor applied to health care, which we call a “health care eBay.” On the health care eBay, four main groups of people “meet” virtually: the patients, the local allied care providers, the remote volunteer doctors, and the medical companies that may provide medicines and supplies.

At least in the beginning stages, we envision the health care eBay to be more of an information delivery mechanism than a remote diagnosis system. A large fraction of the health care problems in rural areas, such as those related to malnutrition, vaccination, access to clean water, care of infants, can be prevented if the villagers had proper access to relevant information. The problem with traditional information delivery mechanisms, such as pamphlets, is lack of means of interaction and lack of pinpoint relevance: one is more likely to ask a question that addresses an immediate and specific concern or need than is to read a generic pamphlet.

In the health care eBay, villagers go to a nearby village clinic (which is manned by a staff who operates the equipment, and is analogous to the village schools discussed in previous sections) to seek health care advice. Digitized voice can be helpful for recording conversations (especially with illiterate patients). Digital images can be useful for documenting externally visible symptoms. Digital video can be useful for documenting motion-related symptoms. Other types of inexpensive and simple diagnostic equipment may also be used. The staff performs data entry and “uploads” all the patient information into the repository in a way analogous to how homework submissions are made (described in Section 8).

Qualified medical professionals, including those overseas, volunteer their time and gain access to the repository in a way similar to how homework graders use the repository (also described in Section 8). (Translation of the recorded voice content can be done by third party translators via the repository as well.) The medical professionals may browse, search, or prioritize the cases that they can help with. They may provide advices, refer cases to other doctors, request more information, or provide more specific instructions to the local staff. Of course, in some cases, the patients must be instructed to go to a

more conventional hospital to seek further diagnosis or treatment.

We may complement the high-bandwidth but asynchronous Postmanet communication mechanism with a low-latency synchronous communication link, such as a cellular phone, or one built on top of packet radio. One possibility is to use the Postmanet to upload the initial bulk data into the repository, establish contact with a remote medical professional who is willing to examine the case, and make an “appointment.” At the appointed hour, the villager and the remote doctor may then communicate with each other using a synchronous voice communication channel. Alternatively, a patient may seek initial diagnosis or treatment at a conventional hospital, but communicate with care providers on the health care eBay for followups.

Although we primarily envision the professionals to be volunteers, we do not rule out for-profit models, and there may be multiple ways of how the medical professionals may collect some fee. For example, donors may “play” on the site as well. Similarly, medical companies may donate (or provide at a low cost) medicines and supplies to the site, some of which may be shipped to the village clinics, where the local staff is responsible for dispensing them.

Of course, stringent security, privacy, and anti-fraud mechanisms must be built into such a system. For example, patients must obtain services from the system only through the screened and trained local staff, with whom we establish a degree of trust, so, for instance, it would be harder for fraudulent patients to fake illnesses and steal free medicines. And only reputable doctors and medicine suppliers are allowed to provide services.

Another way the system can help is to disseminate medical information, such as medical journals and drug information (including dosage information, side effects, and drug interactions) to front-line health care providers. Such information is updated frequently as new drugs and new discoveries are made available. Timely access to such information is badly needed by front-line health care professionals.

11.1.2 What’s New?

Tele-medicine or remote diagnosis in rural areas, such as utilizing high-resolution digital cameras and satellite links to perform retina scans [16], has begun to show promise in recent years. Doctors in the U.S. are also increasingly utilizing emails to communicate

with their patients [17]. These trends, to a certain extent, help us believe that the health care eBay approach discussed here might indeed work. In addition to the potential shown in previous exercises, the health care eBay has several unique features.

- *A pervasive communication mechanism.* The digital communication provided by the Postmanet and phttp, upon which the rest of the system is built, allows the system to reach practically everyone everywhere. This is an advantage that conventional connectivity technologies are unlikely to be able to match any time soon.
- *A cheap and high-bandwidth communication mechanism.* The Postmanet and phttp mechanisms allow us to transmit practically infinite amount of image, voice, video, and other types of medical data virtually for free. This is another advantage that other connectivity options cannot provide.
- *A globally accessible “clearinghouse” that matches “supplies” and “demands” for health care services.* Such an eBay-like clearinghouse for health care does not exist today. A place that aggregates supplies and demands for health care services across time and space can allow doctors to contribute in a manner that is flexible in terms of their time and location commitments. It gives a whole new meaning to the phrase “doctors without borders.” It also helps increase specialization and efficiency.

The health care eBay and the education system (or the learning eBay) that we are working on share a great deal of synergy. The software and hardware infrastructures that we have discussed in previous sections, such as the robotic arm-operated DVD processor, crucial for the automation of the repository operations, can be leveraged by both applications. Indeed, it is conceivable that in some locations, the village school and the village clinic could be co-located.

11.1.3 Potential Benefits

The health care eBay may deliver a number of potential benefits. Some of these benefits are shared by any remote diagnosis systems, while the other benefits are unique or more pronounced in our proposed system.

One of the most important functions of the system is to allow competent medical professionals to contribute at a time and from a location of their choosing. Good doctors take much time and resource

to train and their services are always in great demand everywhere. Given options of where to practice, many doctors are reluctant to serve in rural areas or in the developing world. India, for example, experiences a serious “doctor drain” today: a large number of competent doctors practice in western countries, where they earn much more than at home.

Having made the move, however, many still feel ambivalent about their choices; and if given an alternative which allows them to serve those left behind in some capacity, many would be eager to embrace the opportunity. Similarly, many doctors who live in urban areas would be willing to serve a rural clientele if flexible arrangements can be made. While any remote diagnosis system may help and there are other venues for people to volunteer their services, the health care eBay, as a central clearinghouse, may be particularly effective at harvesting and channeling this pool of volunteerism. One may, for example, easily choose to answer just a few medical email questions about a particular medical issue from a chosen region per week—such “fine-grained” small contributions could have easily been lost without this system, but when successfully harvested, these small efforts from many may add up to a large overall impact.

A browse and search interface of the clearinghouse allows doctors to efficiently and flexibly choose what cases they desire to examine. They may group similar cases together so they can all be dealt with in a similar fashion. They may prioritize based on urgency, so the more urgent cases are less likely to be lost in the “clutter” of a system. They may choose specific regions or communities that have the greatest needs to serve first. All these factors may allow doctors and potential patients to be more efficiently connected with each other, and allow doctors to better utilize their time. (While maximizing efficiency is important, the system does not rule out more personal connections so, for example, it is possible for a doctor and a remote patient to maintain a longer-term connection.)

In traditional settings, where the nearest hospital may be many miles away, transportation is difficult or costly, and a patient is often made to stand in long queues for hours, the patient becomes more reluctant to seek early diagnosis or treatment. Valuable time can be lost, and by the time the symptoms worsen to the point that the patient is forced to seek care, interventions can be too late, too expensive, or too difficult to make a difference. A benefit of any remote diagnosis system is that by making the health care

information and services more easily accessible, patients in remote regions are likely to seek care earlier and more often in the early stages of the symptoms. Such early interventions have been shown to significantly reduce the cost of the health care system.

The case of early intervention can be especially strong for the system that we propose. Because of the pervasiveness and the low cost of the Postmanet communication mechanisms, the infrastructural needs of our proposed health care eBay system are minimal. The low cost and ease of deployment of such a system thus can make it more widely available than other remote diagnosis systems. A greater number of villages may be reached; the overhead for patients to seek care at these village clinics can be even lower; so patients may be even more likely to seek early care.

11.2 Voice Mail

The poor spends disproportionately more (in terms of time, labor, and money) on communication [14]. Previous ICT experiences in the developing countries [3] suggest that variations of voice or email communication are among the most useful and most popular “killer apps.” Success stories include the well known Grameen “phone ladies” in Bangladesh (who loaned cell phones to fellow villagers for a fee), the Brazilian “virtual telephones” (which provided voice mail boxes to those who have no regular telephone service), and a South African experiment that provided email service to residents at post offices.

In this subsection, we discuss a voice mail application built on top of all three key infrastructural components: the Postmanet, EdTV, and the repository abstraction. Although it is not necessary to have all these components, a voice mail system built on top of all three can offer some unique advantages: the Postmanet can offer large bandwidth and universal access at a low cost, EdTV plays the role of demand aggregation and makes the service available to anyone in front of a conventional TV, and leveraging the repository abstraction makes the application easy to build. The voice mail application is not only useful as a standalone service, but also as a part of other higher-level services, such as the distance learning and health care applications that we have discussed in previous sections.

The voice mail application uses the EdTV setup shown in Figure 4. An end user in front of a regular TV screen “activates” the voice mail application using an inexpensive input device (which is either a

walkie talkie or a custom-made ham radio transmitter). The voice mail application runs on the village computer, which receives the end user input via a radio receiver directly attached to it. Once the application is instructed to enter the “record” mode by the end user, it can simply capture the user voice input verbatim. To simplify input tasks such as specifying destination addresses, we can ask users to “hardwire” a number of possible recipient addresses ahead of time (as they “sign up” for the service, for example, potentially in front of the village computer, with the aid of a staff). While the availability of a TV screen is not strictly necessary for this setup to work, it can significantly simplify the user interface.

The captured voice mail may leave the village computer via the Postmanet on a DVD. This arrangement allows a user to speak as much as she desires virtually for free! At a place where other types of connectivity alternatives are non-existent, the Postmanet is the only option. At a place where a low-latency low-bandwidth alternative, such as a cellular link or a packet radio link, exists, a user of the voice mail system may choose between this more expensive, lower-bandwidth, but instantaneous alternative and the cheaper, larger-capacity, but slower Postmanet. The incoming voice mails are delivered over the TV signal sent by the transmitter attached to the village computer.

Some general issues of EdTV that we have discussed in Section 7.3 obviously apply to the voice mail application: one may get a “busy signal” if another person in the same village is already using the system; and there is no privacy as everyone within the transmission range can hear or see the radio and TV signals—if privacy is a concern, one has to walk to the village computer and use the voice mail application there. As we have discussed, we may pick and choose which of the three infrastructural components we desire to use: we may bypass the EdTV component and gain some privacy at the expense of losing some convenience, or we may bypass the Postmanet component and gain some latency advantage (if a cellular or packet radio link is available) at the expense of limited capacity and higher cost. And the loss of privacy over EdTV can also be turned into an advantage: for example, voice mails broadcast or collected over the air wave allow group communication opportunities that one may not usually have easy access to.

11.3 Shopping on EdTV

In this subsection, we discuss a hypothetical shopping application. The purpose of this discussion is not necessarily about this particular application scenario per se—instead, we use this application as a representative of a class of multi-user services that can be delivered over EdTV.

11.3.1 How Does It Work?

We again assume the EdTV setup of Figure 4. Unlike in the voice mail application discussed above, we now assume the presence of a human operator manning the village computer when the shopping service runs. At an hour pre-agreed upon, villagers tune their TVs to a shopping channel, whose signal is fed by the transmitter attached to the village computer. The village computer operator browses a shopping web site at a leisurely pace. Whatever the operator sees on the village computer screen is transmitted to all the TV screens in the village.

If the village computer has a conventional Internet connection, the content of the shopping web site may be retrieved from a remote web server in real time. If the village does not have a conventional network link to the outside world, the content of the shopping web site could have been delivered to the village computer earlier via the Postmanet. (See Section 5.)

As the villagers watch the shopping site on their TV screens, they simply jot down their desired transactions using pencils and paper. (Note that the “shopping site” here can also be an auction site.) Using walkie talkies, the villagers may send the operator some simple requests such as pausing on a page or going back to a previous page. If there are conflicting demands, the human operator plays the role of an arbitrator (like a DJ).

Note that the operator essentially becomes an integral part of a multi-user interface for a shared shopping application in this scenario. An existing traditional shopping site may not even need to be modified. Although it is theoretically possible to build a more automated user interface that eliminates the need of the human operator and to allow villagers to directly control the village computer using their inexpensive input devices (discussed in Section 7.4), this degree of automation may not be necessary and the human operator solution has its advantages that may be difficult to replicate.

The next day, the villagers walk to the “village TV station” and hand their pieces of paper, which contain

their desired transactions, to the village computer operator, who enters the villagers' desired transactions into the computer. These electronic requests would leave the village computer via a conventional network connection, if one is available, or otherwise, a Postmanet connection.

Depending on the nature of the transactions (such as whether it is an auction), the outcomes of the buy and sell "orders" placed by the villagers in earlier days may be electronically delivered to the village computer via some sort of network, and further EdTV viewing sessions, which can be either personal or shared, may be scheduled by the operator to deliver the "news" to his fellow villagers.

Note that all we have accomplished at this point is "information hookup"—the exchanges of physical goods yet need to take place. The information hookup, by itself, however, is valuable. For example, in a traditional setting, a villager may produce some goods on a hunch, walk many miles to a bazaar, and wait hours for a potential buyer to stumble upon his goods, if he is lucky. With an information hookup, the producer may expect to reach a wider audience, and only does the walking (or even the production) when he knows that a pre-arranged buyer is waiting for the goods being sold. This is a more efficient process. More efficient physical exchanges may also be possible when the information hookup permits better demand and supply aggregation, so for example, an enterprising third party could drive an auto-rickshaw to the bazaar on behalf of the whole village on a weekly or daily basis to handle the physical exchanges pre-agreed upon electronically.

11.3.2 Potential Advantages

One may wonder what this shopping application provides that a shopping channel on nation- or state-run TV cannot do. The differences are big. First, as we have discussed in Section 7.2, EdTV is *not* TV: a regular TV is a mass media device, while an EdTV is a more personal media device that permits a high degree of customization. In the shopping example, the goods that are being transacted and the times at which the shopping channel is played can be highly customized for the regions covered by small EdTV transmission ranges. Second, the EdTV shopping service can be "powered by" a site that dynamically pools requests from anyone who has any type of connectivity. This is a level of interactivity that regular TVs do not typically provide. Third, the shopping

mechanism discussed above requires almost no infrastructural support: the villagers do not need access to a single phone, and the village computer powering the transmitter does not even necessarily need a conventional network connection. Yet, the combination of the Postmanet and EdTV allows us to build a fairly sophisticated shared application, combining interesting features of existing Internet shopping and auction sites with that of existing TV shopping channels.

The service described here also has potential advantages over a shopping application delivered via a traditional kiosk. First, an important function of EdTV is to place interesting Internet applications on TV screens inside people's homes, where end users may feel more comfortable spending a larger amount of time passively viewing it, leaving the "driving" to the village operator, and not having to worry about tying up a public device for long. A kiosk in a public location is a less comfortable way of "watching TV," which may explain why, while there are public phones, "public TVs" are not exactly popular. Second, unlike a conventional kiosk, which may incur a high cost on its network connection, the village computer in the proposed system, can get a high-bandwidth feed via Postmanet virtually for free. The low cost and an improved experience (due to the high bandwidth) may attract more usage. The fact that most aspects of a shopping application are intrinsically asynchronous makes it ideally suited for a high-latency network like the Postmanet.

Having speculated about the potential advantages of such a hypothetical application, however, we recognize that realistically building such a service may involve many other complications, and we can be by no means sure that it would work. As we have said, the purpose of this discussion is not necessarily about this application per se, and its details are not important. Our purpose is to use this example service to illustrate how a larger *class* of shared applications may work on top of the infrastructural components discussed in this paper.

12 Conclusion

The Digital StudyHall system includes the following key components. (1) The repository is the centerpiece. All the education content is collected in the repository, and all the digitally enabled human-to-human interactions are carried out through the site. For example, the lecture capture, homework feed-

back, and question-answer session workflows are funneled through the repository abstraction. (2) The Postmanet, and a particular application of it, the phttp mechanism, make the repository available to fringe villages. The Postmanet provides pervasive, high-bandwidth, and low-cost asynchronous connectivity to just about any place, including the remotest areas. (3) EdTV allows us to turn regular TV screens into networked thin client displays. It lowers the cost of end user devices, and it truly bridges the last mile by leveraging TV signals in one direction and radio control signals in the other. EdTV opens up a “portal” for interesting computer-based applications into a poor household that would have otherwise been beyond the reach of more conventional technologies.

These components share a great deal of synergy. The phttp-accessible repository provides an abstraction that is akin to that of a distributed file system, even to places that have no conventional connectivity. The combination of the Postmanet and EdTV provides a natural two-hop solution. And the repository abstraction provides a convenient abstraction to build shared EdTV applications with. A common theme of these components is that they allow sharing and the delivery of highly customized content and experiences.

Using these components, we are building a Digital StudyHall that allows resource-starved village schools to benefit from the better human and content resources available in the urban environments. We hope to connect learners and teachers across time and space, and allow volunteers all over the world to contribute toward the Millennium Development Goal of providing quality universal education!

References

- [1] R. Anderson, M. Dickey, and H. Perkins. Experiences with tutored video instruction for introductory programming courses. In *Proc. of the Thirty-Second SIGCSE Technical Symposium on Computer Science Education*, 2001.
- [2] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 Inside-Out. In *Second Workshop on Hot Topics in Networks (HotNets-II)*, November 2003.
- [3] G. Caspary and D. O’Connor. Providing Low-Cost Information Technology Access to Rural Communities in Developing Countries: What Works? What Pays? (Working Papers No. 229). In *OECD Development Center, Working Papers*, November 2003.
- [4] Committee on Learning Research and Educational Practice, Commission on Behavioral and Social Sciences and Education, and National Research Council, Edited by M. S. Donovan, J. D. Bransford, and J. W. Pellegrino. *How People Learn: Bridging Research and Practice*. National Academy Press, Washington, D.C., 1999.
- [5] J. Dutra, J. F. Gibbons, R. L. Pannoni, M. J. Sipusic, R. B. Smith, and W. R. Sutherland. Virtual Collaborative Learning: A Comparison between Face-to-Face Tutored Video Instruction (TVI) and Distributed Tutored Video Instruction (DTVI). Technical Report TR-99-72, Sun Microsystems Laboratories, January 1999.
- [6] J. F. Gibbons, W. R. Kincheloe, and K. S. Down. Tutored Videotape Instruction: A New Use of Electronics Media in Education. *Science*, 195(3), 1977.
- [7] J. Gray and D. Patterson. A Conversation with Jim Gray. *ACM Queue*, 1(4), June 2003.
- [8] Indian Express, The. EDUSAT is up and ready; not so on the ground. In *The Indian Express*, February 14 2005.
- [9] International Telecommunications Union. Statistics: mobile cellular subscribers. http://www.itu.int/ITU-D/ict/statistics/at_glance/cellular03.pdf, 2003.
- [10] A. Krishnakumar. The decline of public education. *Frontline*, 21(16), July 2004.
- [11] J. Lai, E. Ziskind, F. Zheng, Y. Shao, C. Zhang, M. Zhang, N. Garg, S. Sobti, R. Wang, and A. Krishnamurthy. Distance learning technologies for basic education in disadvantaged areas. *The Eighth Global Chinese Conference on Computers in Education*, June 2004.
- [12] E. Lazowska. UW CSE Tutored Video Pilot Project. <http://www.cs.washington.edu/education/TVI/gibbons.html>.
- [13] MacWorld. Sony is developing 200GB Blu-ray storage. <http://www.macworld.com/news/2004/09/21/-sony/index.php>.
- [14] K. S. McNamara. Information and Communication Technologies, Poverty and Development: Learning from Experience. In *A Background Paper for the infoDev Annual Symposium, The World Bank*, December 2003.
- [15] Ministry of Communications, Government of India. Minister’s address to the Parliament. <http://164.100.24.208/lsq/quest.asp?qref=6546>.
- [16] New York Times, The. For an Ailing Retina, Instant Diagnosis From Afar. In *The New York Times*, July 2 2003.
- [17] New York Times, The. Digital Rx: Take Two Aspirins and E-Mail Me in the Morning. In *The New York Times*, March 2 2005.
- [18] Science Blog: UC Berkeley alumna receives top award for educational reforms in her homeland of India. <http://www.scienceblog.com/community/older/archives/G/ucb133.html>.
- [19] Scientific American. More Bits in Pits. In *Scientific American*, February 2005.
- [20] R. Y. Wang, S. Sobti, N. Garg, E. Ziskind, J. Lai, and A. Krishnamurthy. Turning the Postal System into a Generic Digital Communication Mechanism. In *Proc. ACM SIGCOMM 2004*, August 2004.