

Tangible User Interfaces, a Brief Walkthrough

TUIs have been around quite some time. It could be said that they came well before the GUI as it might be hard to argue against the tangible interaction of a knob or slider moved on the large mainframes of the early computing age. Largely the CLI (Command Line Interface) and GUI (Graphic User Interface) were logical and necessary improvements over the teletype, punch card, and other physical methods used to program and communicate with computers. They certainly lowered the barrier into computing for the general public and it is hard to imagine a computer today that isn't paired with a GUI. The coming decade might prove to be a formative shift in interaction with computers considering the recent trends in VR and Holography for data visualization and the Internet of Things (IoT) and Ambient Intelligence (Aml) making its way into people's houses as an alternative interaction methodology.

So where do TUIs fall in this future? Well that's where things get a bit hazy. Certainly they fall into ubiquitous computing same as IoT or Aml, though as this paper will explain, they also form a complex interaction with data visualization. It may be more appropriate to think of them as a bridge between these paradigms serving both as input and output mediums. First we need to build our vocabulary and understanding of TUIs, following that we will examine some examples of TUIs and discuss if they satisfy our understanding of what TUIs should be, then finish up with some research examining the effectiveness of TUIs.

At this point you probably have some assumptions about TUIs just from the acronym, most likely they may not be as far off from what a TUI is. However, in this section we will explore papers by Hiroshi Ishii, a pioneer in the field and head of the Tangible Media Group at MIT, and Jörn Hurtienne, a professor of psychological ergonomics at the University of Würzburg, in order to really understand TUIs and types of interactions they elicit.

In his 2008 paper *Tangible Bits: Beyond Pixels*, Ishii lays down some fundamental aspects of what a TUI should be. Before that however, he explains GUIs briefly and it's important that we have this working knowledge to inform us about what certain elements belong where and how

they are conceptually different. He says the GUIs are malleable, maybe not in the ways that we imagine something as being physically malleable but they can perform many different actions and operations because the pixels we see on screen can represent data in many different forms. Ishii's argument though about the current model is that generic input methods we currently use do not take advantage of the dexterity our motor skills afford us as humans and for the most part aren't exactly informative of the actions we make. The mouse and keyboard afford input into as system but in terms of output are quite limited beyond the tactile sense we have of physically touching the key and moving the mouse.

That is rule one of Ishii's TUI, it must afford input and output that he labels as control and representation, which is probably a better term when we talk about TUIs. A component of a TUI must act as a control for the data and must also represent the data in some fashion. He further breaks down the TUI as having multimodal representations i.e. the tangible object itself and the intangible data representation. He states that these representations need to be coupled together as one. Ishii explains this in a straightforward manner since physical objects often aren't malleable, not in the way GUIs are as discussed earlier, that they need some augmentation so that a user knows they are manipulating data. This is usually represented by visual or audio feedback though there is also physical responses an object can have using haptics or stimulate our other senses using olfactory and gustatory responses.

Another rule of Ishii's is that objects should act as we expect, if an object takes the form of something it should have the same manipulatable properties that object would have if it was being used normally. He gives the example of a bottle having a cork removed from it, which admittedly is a strange example to consider in terms of how it would be used as a TUI, but for the sake of argument it is an interaction one can imagine in their mind. Perhaps a simpler and more intuitive example is that the object is a knob, it should react to rotation. Things may not necessarily have to respond to six axis manipulation but its manipulation should be something readily understandable by a user and be hinted at by its form.

Ishii next rule is the idea the interaction needs to be seamless. While it is not something unique when we consider the gulf of execution and expectations in a software environment this means

that so too are TUIs tied to these gulfs. Ishii does mention that TUIs have an advantage of sorts though because there is the immediate feedback people receive from the passive tangible object. People can manipulate it and instantly know they have manipulated the physical object, though obviously if the digital “intangible” representation is not in sync the user may experience the same frustrations that would be present in the software environment.

Finally, though not listed formally as one of his rules, more as a design suggestion, TUIs are and should be extremely domain specific. He says that it is important for each tangible tool to be created so that it indicates the function it affords. TUIs are, as he calls it, space-multiplexed input machines, meaning that not only can a user use multiple controls at once, so too can multiple users interact simultaneously. That is what it is important to give specific forms to each controller so that users know what each does in contrast to the others. There is some overlap of these paradigms but some nuances between them too which are useful for understanding the topic. Ishii also gives some examples of TUIs but for now let’s move on to increase our design vocabulary with Hurtienne’s paper.

Unlike Ishii’s paper, Hurtienne uses empirical evidence to determine aspects of TUI design for the tangible components of the interface. His idea was that TUI objects could be built upon population stereotypes that he gathered from linguist’s classifications of conceptual metaphors and image schemas. His research only covered a small portion of these mappings but the findings were significant in their possible design implementations that they put forward. For the sake of ease I am including the tables from his study here which will then be discussed based on his own discussion and also on knowledge of perception and cognition of humans.

Mapping							Overall	
	%	str	%	str	%	str	%	str
MORE IS BIG – LESS IS SMALL	100	1.00	100	1.00	100	1.00	100	1.00
POWERFUL IS BIG – POWERLESS IS SMALL	100	1.00	93	0.87	93	0.86	96	0.91
IMPORTANT IS BIG – UNIMPORTANT IS SMALL	79	0.57	87	0.73	88	0.75	84	0.69
SIGNIFICANT IS BIG – INSIGNIFICANT IS SMALL	87	0.73	80	0.60	73	0.47	80	0.60
KNOWING IS BIG – UNKNOWING IS SMALL	53	0.07	86	0.71	80	0.60	73	0.45
VALUABLE IS BIG – VALUELESS IS SMALL	53	0.07	71	0.43	87	0.73	70	0.41

Table 1. Results of BIG-SMALL mappings

Mapping							Overall	
	%	str	%	str	%	str	%	str
GOOD IS BRIGHT – EVIL IS DARK	100	1.00	67	0.33	71	0.43	80	0.59
ANGRY IS DARK	86	0.71	93	0.86	63	0.25	80	0.59
HAPPY IS BRIGHT – SAD IS DARK	87	0.73	73	0.47	73	0.47	78	0.56
GOOD IS BRIGHT – BAD IS DARK	93	0.86	44	-0.13	67	0.33	67	0.33
RELIGIOUS IS BRIGHT – GODLESS IS DARK	80	0.60	87	0.73	29	-0.43	66	0.32
INTELLIGENT IS BRIGHT – STUPID IS DARK	64	0.29	25	-0.50	40	-0.20	42	-0.16

Table 2. Results of BRIGHT-DARK mappings

Mapping					Overall	
	%	str	%	str	%	str
EMOTIONAL IS WARM – UNEMOTIONAL IS COLD	91	0.83	95	0.91	93	0.87
ANGRY IS WARM	83	0.65	82	0.64	82	0.64
HAPPY IS WARM – SAD IS COLD	65	0.3	91	0.82	78	0.56
PROBLEMATIC IS WARM – UNPROBLEMATIC IS COLD	68	0.36	87	0.74	78	0.56
FEARFUL IS COLD	68	0.36	57	0.13	62	0.24
DISGUSTING IS WARM	70	0.39	50	0	60	0.2

Table 3. Results of WARM - COLD mappings

Mapping		
	%	str
MORE IS HEAVY – LESS IS LIGHT	98	0.96
IMPORTANT IS HEAVY – UNIMPORTANT IS LIGHT	93	0.87
SAD IS HEAVY – HAPPY IS LIGHT	89	0.78
PROBLEMATIC IS HEAVY – UNPROBLEMATIC IS LIGHT	84	0.69
EXPENSIVE IS HEAVY – CHEAP IS LIGHT	80	0.60
INDEBTED IS HEAVY – DEBT FREE IS LIGHT	75	0.50

Table 4. Results of HEAVY-LIGHT mappings

Mapping		
	%	str
BORING IS SMOOTH – DANGEROUS IS ROUGH	84	0.69
POLITE IS SMOOTH – IMPOLITE IS ROUGH	82	0.64
CONTINUOUS IS SMOOTH – INTERRUPTED IS ROUGH	82	0.64
UNPROBLEMATIC IS SMOOTH – PROBLEMATIC IS ROUGH	82	0.64
PLEASANT IS SMOOTH – UNPLEASANT IS ROUGH	71	0.42

Table 5. Results of SMOOTH-ROUGH mappings

To quickly summarize what the tables mean, boxes in white are connections people made that were significant and the grey boxes ones that did not meet his cutoff strength value of greater than .60. It should be noted that these mappings were tested on German residents, though they were taken from English metaphors and provided by the British National Corpus, these mappings may not translate 100% to American audiences. Still the findings do provide a nice framework for understanding how the interactive objects of a TUI can utilize their form, color, weight, texture, and temperature to inform a user about its “intangible” data representation which Ishii outlined in his paper. This gives designers a lot more tools in which to design the interfaces with considering that some of paradigms and designs for GUIs that have been developed over the last 40 years may not immediately apply to TUIs.

There is also some overlap though, especially where perception is common between the interfaces. Designers can still rely on certain colors and their meanings, however, because of

the tangibility of the objects, designers should be aware of other meanings as well. While Hurtienne was looking for the difference between light and dark with those metaphorical markings (table 2) there was a lot of different associations with color and shades of those colors that made such mappings far more complicated. Perhaps most illustrative of this came during his discussion of the results is his remark “The light blue brick was sometimes seen as bad, it had associations with mold, being pale and unhealthy, but then again associated with the sky and happiness.” (Hurtienne, 66) Still, some stereotypes were quite strong like red meaning hot and blue meaning cold, we can probably still rely on certain perceptual concepts created for GUIs that would inform users of TUI interactions like a red warning label. Certainly those mappings held over and certainly other mappings that work in a traditional GUI environment should also be applicable to TUI environments. The benefit of this study is that it shows that metaphoric representation can be used with physical components of TUIs to fill out interactions and indicate responsiveness to users in multiple dimensions.

So now that we have an understanding and vocabulary of TUI's we can examine some that succeed, fail, and maybe land somewhere in-between. First up is the TUI d-touch which was created and studied by Enrico Costanza. The d-touch is a simple, low cost TUI that allows users to make beats using two modes utilizing the concepts of an audio sequencer and a drum machine. The purpose of the study was to examine effectiveness and responses to a TUI that could be made with low cost parts like a webcam, a personal computer, and a printer. The system was calibrated and controlled using camera vision and software that could be ran in a browser. The benefits of the system was it was easily scalable and configurable utilizing AR (Augmented Reality) recognition to handle the initial setup. All one had to do was print out the fiducial markers and the sheet and place it underneath their webcam. Certain fiducials did certain functions and because AR can determine rotation from the fiducials, interaction could be mapped to that aspect as well the markers position. The study was considered a moderate success though users didn't continue after the initial novelty wore off as they discovered the average user logged about 5.55 sessions for 7.34 minutes. Few users collectively put more than 90 minutes into the interface but considering its limitations it should not be held against it.

Using our vocabulary from before we should consider if it works well as a TUI or not. It passes the first test as it enables input and we are given audio output in response. It also passes the responsiveness test as their study didn't indicate that it was a major problem, though they do mention that some users expressed issues with recognition and that perhaps lighting played a role. Still responsiveness of the environment can be considered a pass. That is about all though, the issue with d-touch is that it's overly generic, the markers don't inform the user as to their actual function so that they are not domain specific. Users can inference some functionality from the drum machine since things are labeled but fiducials don't particularly tell the user anything about what actions or functions they perform, not without some visual augmentation. This environment must be explained, which isn't entirely problematic since all environments require some amount of explanation, but it doesn't do anything to teach its users with the physical objects thus ignoring one of TUIs greatest strengths. The idea is commendable since it offers a low cost alternative to TUIs but it also isn't very domain specific and is also an adaptation of a field which has many TUI's that already exist. Audio equipment, probably more than any other modality, has always had TUI's. Before we could make audio using GUIs people used hardware synthesizers that could easily be considered early TUIs. Even the functions d-touch offers are pretty standard audio hardware functions of a drum pad or sequence synthesizer. A simple search on Amazon.com can find equipment that serves the same purpose for low costs and is more reliable. Thus, while the novelty was interesting, it may not have been a very successful TUI.

The next TUI is Florian 'Floyd' Mueller's Air Hockey Over a Distance. This TUI is pretty straightforward. For anyone who has never experienced an air hockey table it is a perforated table that blows air out in order to make a mostly frictionless plane for a plastic puck to float on that you hit back and forth attempting to score it into your opponent's goal. What Mueller did was essentially bisect a table and allowed users to play the game physically with others over a network. It used all the hardware of a normal table plus a projection screen, some sensors, and components in which to fire the puck back at users. In this first model the puck would fire randomly, not based on the trajectory of what the opponent hit back, but it was admittedly a time constraint the developers ran into. They mentioned some other iterations and examples

that used a virtualized puck that was projected but also that users enjoyed the system they had even if the puck didn't line up exactly.

So how does this compare against Ishii's definition. There certainly appears to be input and output though it is a bit more complex of an interaction. The puck itself and the mallets don't map to any direct input but through them you interact with

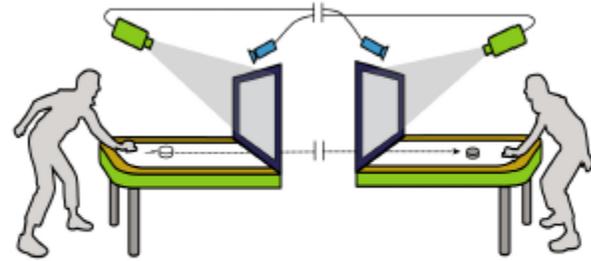


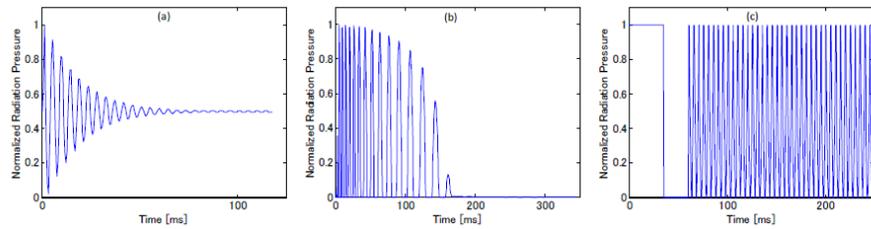
Figure 2. Conceptual design.

the system and your opponent. The interaction can be seen on the screen and is also exemplified when the puck is fired back at your side. The interaction is pretty seamless and interacts mostly how we would expect. The system uses the same objects you would normally use to play air hockey so the objects tell you how to use them, the only issue is that the puck does not have the 1:1 ratio which can affect the seamlessness of actually competing against an opponent. This was mentioned in their report that users had wished the puck mapping was 1:1. The system is also incredibly domain specific so it appears to pass all of Ishii's rules. The mention of a virtual puck would have probably fixed the seamlessness as it is easier to map the pucks movements between locations that way, however, it would lose a lot of the physicality that makes it so charming. The mallets then would have to be outfitted with some form of haptic feedback and audio cues added or else the illusion is broken that you are actually playing air hockey. Using these virtual representations would present a more challenging task requiring that the system acts in the same way that the physical analog world does, or how we expect it to.

The final TUI examined is Yasuaki Monnai's HaptoMime. This TUI will probably see use for things which have projected images though need some form of tactile input. The system they created uses ultrasonic acoustic radiation pressure to simulate tactile feedback to a user. They use a 40 kHz signal that can be modulated with other frequencies to give different forms of feedback which they describe as being stiff and light, a burst of air, and vibrating. For their specific application they created it to give feedback to floating images created by using a special mirror which in their application was called an Aerial Imaging Plate (AIP). The problem they

discussed with floating images was that the lack of tactile feedback essentially created a gulf of expectations, it also

created errors and breaks seamless with the interface as the user's finger passes



through the floating image. Their solution however shouldn't be considered as specific to their application, as we work with implementing VR, AR, and holography we need a way to get tactile feedback with the objects we interact with. Take for example a laser keyboard, an interesting novelty that felt promising but without tactile feedback was doomed to user error. Tapping your fingers against a flat surface didn't afford the same type of feedback as the keys on a



keyboard so even if the keyboard was 100% accurate, it was hard to tell if your fingers were resting in the right areas and ultimately resulted in gulfs of execution and expectations.

The technology is certainly interesting but is it a TUI? Maybe. It is a component of the interface though it doesn't take input directly, it must be given a coordinate to give feedback to. In HaptoMime this was an IR sensor and for the laser keyboard it would be the key you are supposed to be pressing. It enables seamless and the illusion that the projected image works as we would expect in terms of feedback. However, HaptoMime essentially has no real tangible objects that afford its interactivity. Rather it is a component of the system and this seems like it could certainly enable feedback but isn't a

TUI in a traditional sense. This technology though could be very useful in future applications when paired with the data visualization methods mentioned at the beginning of this paper.

Enough examples of TUIs, though there is one more question that needs to be answered, are they effective? Initially in his 2007 article *Do Tangible Interfaces Enhance Learning*, Paul Marshall asked that question. He went through and examined a lot of the literature and

projects that had popped up about TUIs and really put forward whether such theories were grounded in any empirical evidence or just were untested assumptions. He certainly leans toward the latter that while some experimentation had been done there was very little work that sought to determine what components of tangible interfaces were supportive and which ones detrimental. He outlined domains with potential for applications of TUIs, noting that the domains that offered the most promise contained abstract spatial components like in molecular chemistry. Other learning activities such as those that elicited exploration or creativity he considered to be prime candidates for TUI integration. Ultimately his goal was to create a framework for discussing the

education applications of TUIs. He created six main elements for this framework to help guide discussion of the education benefits and applications of tangible interfaces to which he argued needed more empirical evidence rather than untested assumptions.

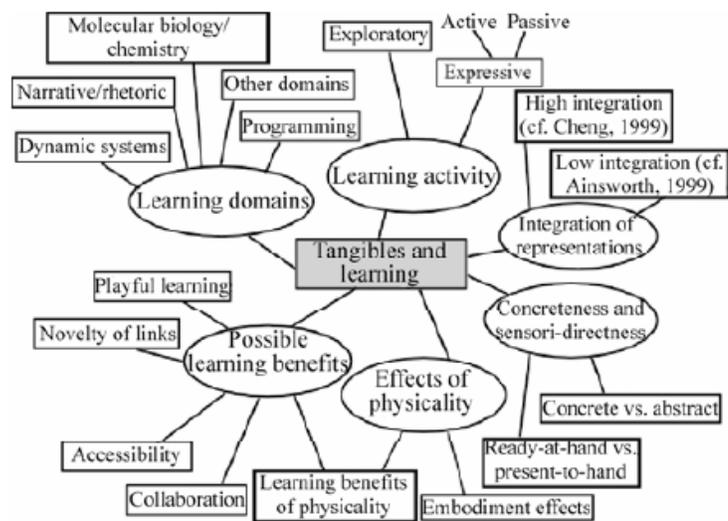


Figure 1: Analytic framework on tangibles for learning

Acknowledging his work was Sébastien Cuendet in his study in 2012. He took Marshall's framework and examined a very particular component within it to see what aspects of the TUI were encouraging learning. Cuendet focused his study on how feedback was given to the user and its effect on learning. He believed that TUIs enabled collaboration, were accessible, enabled playful learning, though also fostered novelty. Like Ishii outlined in his rules of TUIs, Cuendet was even more firm that coupling tangible input with visual output was the cornerstone of high usability in a system. While his stance may have been overly literal and ignorant of other modalities, for his study it was crucially important. What he sought was whether real time feedback, which is something a TUI strives to afford, or delayed feedback was more effective in fostering learning. His study examined carpentry apprentices who were asked to recognize the position of a block and its rotation based on different views of the

object, a face view and a side view. The subjects were tested several times and their improvement measured. The exercises they were asked to perform were also measured, there was a score attributed to accuracy and efficiency of their interactions which served to prevent trial and error effecting the actual learning process. What he ultimately found was that the type of feedback affected learning at different levels. In his trials where the participants got real time feedback there was an overall higher accuracy of right answers but a lower level of improvement. In tests where there was delayed feedback they got less correct answers but showed a significant improvement in their understanding over the trial period. What Cuendet believed was that instant feedback helped at the task level, it helped people make more accurate decisions quicker but ultimately left no time for reflection so it didn't stimulate actual learning, delayed feedback on the other hand, helped at the process level, making for less accurate answers in the short term activities but allowed for greater understanding and time for reflection, ultimately stimulating learning.

In 2013, Amanda Strawhacker furthered the study of TUI's educational effectiveness by testing against different modalities when using the same system. Her study looked at TUIs, GUIs and HUIs (Hybrid User Interfaces) which was just a combining of the TUI and GUI elements. In her study she examined kindergarten student and their ability to program a simple robot using the different modalities. The programming process was very abstracted so that aspects of programming were represented by wooden blocks for the TUI components and those same graphics represented on the wooden blocks were also in the graphical programming environment. She ran her test on both traditional kindergarten students and a Montessori school group. They were tested on tasks of sequence of events and loops plus a culminating task, the former two also had easy and hard variations.

	Easy Sequencing	Hard Sequencing	Easy Repeats	Hard Repeats	Culminating Task
TUI	73.3% (n=15)	73.3% (n=15)	46.7% (n=15)	0% (n=15)	33.3% (n=15)
HUI	66.7% (n=12)	33.3% (n=12)	41.7% (n=12)	0% (n=12)	18.2% (n=11)
GUI	12.5% (n=7)	50% (n=6)	66.7% (n=6)	16.7% (n=6)	16.7% (n=6)

The results were pretty interesting as they were mostly insignificant though each group had its strengths and weaknesses. The TUI group performed the better as an average overall but was

the best in the sequencing components and the culminating task. Meanwhile the GUI group performed the best in the loop components though below average in the other tasks. Meanwhile the HUI group didn't excel at anything and was slightly below average across the board. Strawhacker considered that the students in that group were most likely overwhelmed by the options afforded by dual interfaces. She also believed that opportunities for learning were afforded by the recognizable nature of the tangible components since the kids were familiar with wooden blocks.

While the jury may still be out on whether we can consider if TUI environments encourage learning, it is likely they will play a major role in the foreseeable future due to the IoT and Aml. Technology like that of the HaptoMime may play an integral role in the future of holography and other projected interfaces that would be enhanced by some form of tactile feedback. There is also promise for the future of TUI design based on the research done by Hurtienne where population stereotypes can be built in to TUIs in order to facilitate recognition of components and give them properties that help users understand their interaction as well as understand the visualization of the intangible data. As the medium matures, more and more design paradigms and applications will be built up around new research in the field. The TUIs growth will probably be similar to the growth of the GUI, which was not instantaneous and was built up through empirical studies and design paradigms from reiteration of design over the last 40 years.

Works Cited

1. Costanza, E. Giaccone, M. Kueng, O. Shelley, S. and Huang, J. (2010) *Ubicomp to the masses: a large-scale study of two tangible interfaces for download*. Proceedings of the 12th ACM international conference on Ubiquitous computing (UbiComp '10). ACM, New York, NY, USA, 173-182.
2. Cuendet, S. Jermann, P. and Dillenbourg, P. (2012) *Tangible interfaces: when physical-virtual coupling may be detrimental to learning*. Proceedings of the 26th Annual BCS Interaction Specialist Group Conference on People and Computers (BCS-HCI '12). British Computer Society, Swinton, UK, UK, 49-58.
3. Hurtienne, J. Stößel, C. and Weber, K. (2009) *Sad is heavy and happy is light: population stereotypes of tangible object attributes*. Proceedings of the 3rd International

- Conference on Tangible and Embedded Interaction (TEI '09). ACM, New York, NY, USA, 61-68.
4. Ishii, Hiroshi. (2008) *Tangible bits: beyond pixels*. Proceedings of the 2nd international conference on Tangible and embedded interaction (TEI '08). ACM, New York, NY, USA, xv-xxv
 5. Marshall, Paul. (2007) *Do tangible interfaces enhance learning?* Proceedings of the 1st international conference on Tangible and embedded interaction (TEI '07). ACM, New York, NY, USA, 163-170.
 6. Monnai, Y. Hasegawa, K. Fujiwara, M. Yoshino, K. Inoue, S. and Shinoda, H. (2014) *HaptoMime: mid-air haptic interaction with a floating virtual screen*. Proceedings of the 27th annual ACM symposium on User interface software and technology (UIST '14). ACM, New York, NY, USA, 663-667.
 7. Mueller, F. Cole, L. O'Brien, S. and Walmink, W. (2006) *Airhockey over a distance: a networked physical game to support social interactions*. Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology (ACE '06). ACM, New York, NY, USA, Article 70
 8. Strawhacker, A. Sullivan, A and Umaschi Bers, M. (2013) *TUI, GUI, HUI: is a bimodal interface truly worth the sum of its parts?* Proceedings of the 12th International Conference on Interaction Design and Children (IDC '13). ACM, New York, NY, USA, 309-312.